

Three Grades of Truth Relativity*

John MacFarlane[†]

Draft of October 14, 2003

Please do not quote or circulate without permission.

1 Introduction

Is the *sentence* “You have been to China” true? The question is confused. The sentence can be used to make a true or an untrue assertion, but it is not itself either true or untrue. Hence orthodox truth-conditional semantics assigns it a truth value only relative to a context of use. But what about a particular utterance of “You have been to China”—say, Roland’s assertion on June 14, 2002, in the company of Mia? Here the orthodox view does assign an unrelativized truth value, determined by the sentence and various features of the context of use. Roland’s assertion is true or untrue *simpliciter*; there is no room for *further* relativity once the context of use has been taken into account. In a slogan: sentence truth is relative, but utterance truth is absolute.

I think that the orthodox view is wrong. In this paper, I argue that sentence truth must be relativized not just to contexts of use, but to what I call *contexts of assessment*. If this is right, then we cannot assign unrelativized truth values to assertoric utterances. New slogan: sentence truth is *doubly* relative (to a context of use and context of assessment), while utterance truth is *singly* relative (to a context of assessment). Accordingly,

*This paper descends from talks I gave in May and June, 2002, in Berkeley and Lund, Sweden. A later draft was discussed in September 2002 by the Bay Area Philosophy of Language Discussion Group (BAPHLD), and I am grateful to members of that group, particularly Kent Bach and Ken Taylor, for their comments. For useful discussion and correspondence, I thank Nuel Belnap, Joseph Camp, Jeff King, Peter Hanks, Benj Hellie, Max Kölbel, Berislav Marusic, Elijah Millgram, Ram Neta, Lionel Shapiro, Jason Stanley, and Paul Teller. Finally, I acknowledge the financial support of an ACLS/Andrew W. Mellon Fellowship for Junior Faculty and a Berkeley Humanities Research Fellowship.

[†]Department of Philosophy, University of California, Berkeley. 314 Moses Hall, Berkeley, CA 94720–2390. E-mail: jgm@uclink.berkeley.edu.

we must recognize two very different kinds of linguistic context sensitivity: sensitivity to features of the context of use (*use sensitivity*) and sensitivity to features of the context of assessment (*assessment sensitivity*).

It is easy to confuse assessment sensitivity with something more mundane: the relativization of truth to an “index” (Lewis) or “circumstance of evaluation” (Kaplan). To avoid this confusion, which is rampant in the literature on relative truth, I proceed in stages. Just as Quine famously distinguished “three grades of modal involvement,” I distinguish three grades of truth relativity. The first grade is relativization of truth to a *point of evaluation*; this much relativization is necessary in order to define truth for languages with quantifiers and other non-truth-functional operators. The second grade is relativization of truth to a *context of use*; this kind of relativization is necessary for any language containing tensed sentences, indexicals, or other context-sensitive elements. The third grade is relativization of truth to a *context of assessment*. Unlike the others, this third grade of truth relativity has never been clearly formulated, and consequently it has been little explored.

Of course, it is not enough merely to describe a formal semantic framework that allows the third grade of relativization of truth: one must say what it *means* to say that an utterance is true relative to a context of assessment. I do this by giving an account of assertion that gives a role to assessment-relative truth. This account allows us to say what it would be like to speak a language with assessment-sensitive expressions, while leaving it open whether English is such a language.

I then present what seems to me a compelling example of assessment sensitivity in natural language: future contingents. Suppose I assert “There will be a sea battle tomorrow,” when it is objectively indeterminate whether there will be a sea battle the next day. How should we evaluate my utterance for truth? Given only the first two grades of truth relativity, our options are limited. Utterance truth is absolute, so we must give my assertion an absolute truth value: either true, false, or neither. But there are serious problems with all of these choices. Only by making use of the third grade of truth relativity, I argue, can we open up space for an adequate solution. On this account, my utterance is true as assessed from the future tomorrow with the sea battle, false as assessed from the future tomorrow without the sea battle, and neither true nor false as assessed today. The framework I provide shows how we can make sense of these relativized truth assessments.

Seeing assessment sensitivity at work in this familiar region of natural language should remove any *principled* worries about its coherence—and prepare us to find it elsewhere. In the remainder of the paper, I explore

its applications to David Lewis’s theory of accommodation and to debates about evaluative relativism.

2 Truth relative to a point of evaluation

One of Tarski’s great conceptual insights was the realization that in order to give a systematic account of the way in which the truth values of quantified sentences depend on their construction and the semantic values of their parts, we need to use a *relativized* notion of truth: truth on an assignment of values to the variables, or, as Tarski called it, “satisfaction.” It is this relativized notion, and not truth *simpliciter*, that is defined in the recursive clauses of a Tarskian truth theory. Thus, for example, the clause governing the universal quantifier looks like this:¹

Clause for \forall : ‘ $\forall x_i \Phi$ ’ is true on assignment a iff for every assignment a' that differs from a at most in the value it gives to x_i , Φ is true on a' .

Note that even the clauses for the truth-functional connectives must be stated in terms of truth on an assignment, as these connectives may operate on open formulas:

Clause for \neg : ‘ $\neg \Phi$ ’ is true on assignment a iff Φ is not true on a .

In order to get from truth on an assignment to truth *simpliciter*, we need an additional definition:²

Definition of absolute truth for sentences: If Φ is a sentence, then Φ is true iff Φ is true on every assignment.³

¹I use single quotes as Quinean corner quotes throughout.

²Cf. Herbert B. Enderton, *A Mathematical Introduction to Logic* (Boston: Academic Press, 1972), 84; C. C. Chang and H. J. Keisler, *Model Theory*, third edition (Amsterdam: Elsevier Science, 1990), 32. I assume a fully interpreted language with a fixed domain, so there is no need to relativize truth to an interpretation or a domain. I will discuss relativization to a domain in section 4, below. Relativization to an interpretation does not fall into any of my three “grades:” my project is to classify relativizations of truth for fully interpreted sentences, not for *partially* interpreted sentences containing schematic “nonlogical constants” (see note 20, below).

³Equivalently, “on *some* assignment,” or even “on assignment a_0 ,” since a sentence will have the same truth value on every assignment. Later we will look at structurally similar cases in which the quantification is not vacuous.

This definition removes the relativization to assignments, though only for sentences (formulas with no free variables). Formulas with free variables are not given absolute truth values. It is worth thinking about why. Technically, we *could* extend the definition to open formulas. We would get some semantic anomalies, like true disjunctions without a true disjunct (e.g., “ $Fx \vee \neg Fx$ ”), but if we really *needed* to assign absolute truth values to open formulas, we might just learn to accept these.⁴ Thus the real reason we do not extend the stipulation to open formulas must be that we do not *need* them to have absolute truth values. But why don’t we?

Here is a suggestion. Our practice of assigning truth values to sentences does not stand alone. We care about whether sentences are true because in asserting sentences, we commit ourselves to their truth and represent ourselves as believing that they are true. But we don’t assert open formulas like “ x_1 is a dog.”⁵ (What would we be *saying* thereby?) That is why we don’t need to assign them absolute truth values. We assign them *assignment-relative* truth values only as a technical expedient for systematizing the (absolute) truth values of sentences containing them.

It may be objected that we *do* sometimes assertively utter open formulas, but I think the exceptions prove the rule. Often there is a convention in place of treating assertions of open formulas as assertions of their universal closures.⁶ But in that case what is really asserted is not an open formula at all, but a quantified sentence. Again, one might say: “Let x_1 and x_2 be 17 and 19. Then x_1 and x_2 are twin primes.” Here the variables are being used as temporary names, with particular intended referents, not as genuine (bindable) variables. In every case in which we appear to assert open formulas, there is a special convention in place for treating open formulas as something else.

The first grade of truth relativity is required whenever our language contains non-truth-functional operators. Quantifiers are just one example; modal operators are another. Here we need to relativize truth to *possible worlds*:

⁴Advocates of supervaluational semantics for vague discourse accept similar anomalies: see e.g. Rosanna Keefe, *Theories of Vagueness* (Cambridge: Cambridge University Press, 2000), 181–8.

⁵For this point, see Nuel Belnap, Michael Perloff, and Ming Xu, *Facing the Future: Agents and Choices in Our Indeterministic World* (Oxford: Oxford University Press, 2001), chapter 6, at 157.

⁶Cf. Alfred Tarski, *Introduction to Logic and the Methodology of Deductive Sciences*, second ed., trans. Olaf Helmer (New York: Oxford University Press, 1946), 7–8; Chang and Keisler, *Model Theory*, 32.

Clause for \Box : ‘ $\Box\Phi$ ’ is true at world w iff for every world w' that is accessible from w , Φ is true at w' .

As before, we will need some way to get from the world-relative truth values produced by the recursive clauses of our semantics to absolute truth values. The definition of absolute truth in this case is a bit different: we don’t *quantify* over worlds, as we did over assignments, because the (absolute) truth value of a modal sentence is its truth value at *the actual world*:

Definition of absolute truth for a modal language: Φ is true iff Φ is true at the actual world.

What if our language contains both quantifiers *and* modal operators? Then we will need to relativize truth to both assignments and worlds. It is convenient to think of world and assignment as two *parameters* of a single *point of evaluation*, which we can represent as an ordered pair $\langle w, a \rangle$. Then our recursive clauses will look like this:

Clauses for \forall , \Box , and \neg :

‘ $\forall x_i\Phi$ ’ is true at $\langle w, a \rangle$ iff for every assignment a' that differs from a at most in the value it gives to x_i , Φ is true at $\langle w, a' \rangle$.⁷

‘ $\Box\Phi$ ’ is true at $\langle w, a \rangle$ iff for every world w' that is accessible from w , Φ is true at $\langle w', a \rangle$.

‘ $\neg\Phi$ ’ is true at $\langle w, a \rangle$ iff Φ is not true at $\langle w, a \rangle$.

And absolute truth can be defined as follows:

Definition of absolute truth for a quantified modal language: If Φ is a sentence, then Φ is true iff Φ is true at every point $\langle w, a \rangle$ at which $w =$ the actual world.

As we add more non-truth-functional operators to the language, our points of evaluation will become more and more complex, with more parameters. But no matter how complex the points become, the *purpose* of relativizing truth to points of evaluation is always the same: to facilitate a systematic account of the way in which the truth values of whole sentences depend on the semantic values of their parts and their syntactic mode of composition. Relativization to points has nothing to do specifically with context sensitivity; as we have seen in the quantificational case, it is required even for the semantics of “eternal” sentences like the sentences of mathematics. So even someone who, like Quine, eschews context sensitivity in a canonical language must make use of this first grade of truth relativity.

⁷For simplicity, I assume a fixed and world-independent domain.

3 Truth relative to a context of use

Most sentences in natural languages cannot be assigned absolute truth values. Take “The sun is shining in Berkeley.” This sentence can be used to make a true assertion or a false one, depending on when it is uttered. Thus it does not make sense to say that the sentence itself is true or false, but only that it is true or false relative to a particular *context of use*. Relativization of truth to a context of use may even be needed for a language without tense or indexicals. Contingent but temporally eternal sentences like “Spacetime is curved” can be truly uttered in some possible worlds, but not in others. If we want to track this variation in truth value across counterfactual contexts of use, then we should ask our semantics to deliver context-relativized truth values, rather than absolute ones.⁸

Instead of connecting truth *simpliciter* with truth at a point, we now need to connect truth *at a context of use* with truth at a point. In the simple modal case discussed in the previous section, the task is easy:

Definition of context-relative truth for modal sentences:

Φ is true at context of use C iff Φ is true at the world of C .

Note that the first occurrence of “true at” in this definition refers to truth at a context, while the second occurrence refers to truth at a point of evaluation.

The simplicity of this account might make one wonder why it is so important to distinguish between truth at a point and truth at a context. The reason usually given in the literature for distinguishing points from contexts is that the parameters of points must be individually shiftable in a way that the features of contexts are not.⁹ For example, in order to handle tense operators, which shift the time parameter, we will need to countenance a point of evaluation at which the speaker parameter = Richard Nixon, the location parameter = Moscow, the world parameter = our world, and the time parameter = August 8, 1974. But there is no corresponding context of use, since in our world Nixon was in Washington on August 8, 1974. Relatedly, “I am here” is true at every context of use. But it had better not be true at all points of evaluation, because if it were, “It is always the case

⁸As David Lewis puts the point, “Contingency is a kind of indexicality” (“Index, Context, and Content,” in Lewis’s *Papers in Philosophical Logic* [Cambridge: Cambridge University Press, 1998], 25).

⁹See David Kaplan, “Demonstratives,” in *Themes from Kaplan*, ed. Joseph Almog, John Perry, and Howard Wettstein (Oxford: Oxford University Press, 1989), 507–510; and David Lewis, “Index, Context, and Content,” 28–31.

that I am here” and “It is necessary that I am here” would also be true at every context, and they are not.¹⁰

These arguments certainly do show that not all points of evaluation can be contexts of use. But for all that has been said, we might still identify the contexts of use with a *proper subset* of the points. If for each context of use there is a unique point whose parameters it determines, would it not simplify things to *identify* the contexts of use with these special points? On this view, the set of contexts would be a subset of the set of points, and the relation of truth-at-a-context would simply be a *domain restriction* of the relation of truth-at-a-point.

The reason this proposal won’t work—and a deeper reason for distinguishing between truth-at-a-context and truth-at-a-point—is that in general, there will not *be* a unique point corresponding to a given context of use. To see why, it is only necessary to consider how context-relative truth might be defined in terms of point-relative truth for a language with modal operators *and* quantifiers:

Definition of context-relative truth for quantified modal

sentences: Φ is true at context of use C iff Φ is true at every point $\langle w, a \rangle$ at which $w =$ the world of C .

Utterances take place at worlds, but not at variable assignments. So there will be *infinitely many* points corresponding to a given context of use, all sharing a common world parameter but differing in their assignment parameter. There is no such thing as “*the* point corresponding to the context of use.” Only *some* parameters of points are “initialized” by context in the definition of context-relative truth.¹¹ Those that are not initialized are quantified over. Later, I will give some examples of parameters other than assignments that are not initialized by context.

There is a further difference between contexts and points of evaluation, as I think of them. While contexts of use are concrete utterance episodes (or perhaps episodes of use in thought), points are just abstract sequences of parameters.¹² Points will contain just those parameters that are required in order to give a recursive definition of truth-at-a-point for the language. Contexts, on the other hand, “come with” all the detail one would like. Thus we can talk of the speaker of the context, the audience of the context, the

¹⁰Kaplan, “Demonstratives,” 509.

¹¹The terminology of “initialization” comes from Belnap et al., *Facing the Future*, 148–9.

¹²Here I agree with Lewis: see the Postscript to “General Semantics,” reprinted in volume 1 of Lewis’s *Philosophical Papers* (New York: Oxford University Press, 1983), 230, and “Index, Context, and Content,” 21, 28, 30, 31.

time of the context, the location of the context, the speaker's intentions at the context, the phase of the moon at the context, and so on without limit. Of course, not every feature of context will be semantically relevant. Which features are semantically relevant will depend on the particular expressive resources of the language.

It is worth pointing out two important formal differences between the framework for semantics I am presenting here and the more familiar frameworks of Lewis and Kaplan. First, my *points of evaluation* contain all the parameters to which truth must be relativized in order to get a recursive definition of truth (at a point) for a given language. In the simple language considered above, the points include just two parameters (world and assignment), both of them capable of being “shifted” by some operator in the language. But in more complex languages, the points might include parameters that are not shifted by any operator, like a speaker parameter (for a language containing “I”) or a time-of-utterance parameter (for a language containing “now”). Thus, points of evaluation include both shiftable and non-shiftable parameters.¹³ Lewis and Kaplan, by contrast, make a terminological distinction between shiftable and non-shiftable parameters. The shiftable parameters make up an “index” (Lewis) or “circumstance of evaluation” (Kaplan), while the non-shiftable parameters are provided by the “context.”

As a result, for Lewis and Kaplan, context must play a role in the recursive clauses of the semantics (which need to know about the non-shiftable parameters). Lewis's recursive clauses define “sentence s is true at context c at index i ”,¹⁴ while Kaplan's define “sentence s is true at context c , assignment f , time t , world w .”¹⁵ Context then plays another role in the definition of truth at a context, where it “initializes” the shiftable parameters. For Lewis, “sentence s is true at context c iff s is true at c at the index of the context c .”¹⁶ For Kaplan, formula ϕ is true at context c iff ϕ is true for every assignment at the time and world of c .¹⁷ On the framework advocated here, by contrast, context plays no role in the recursive clauses of the semantics, which simply define “ Φ is true at point p .” The sole role of context is to initialize parameters—both shiftable and non-shiftable—in the definition of truth at a context.

These differences are perhaps merely notational. Certainly nothing of

¹³Here I follow Belnap et al., *Facing the Future*, 142–5.

¹⁴“Index, Context, and Content,” 31.

¹⁵Paraphrasing “Demonstratives,” 544.

¹⁶“Index, Context, and Content,” 31.

¹⁷“Demonstratives,” first definition on 547.

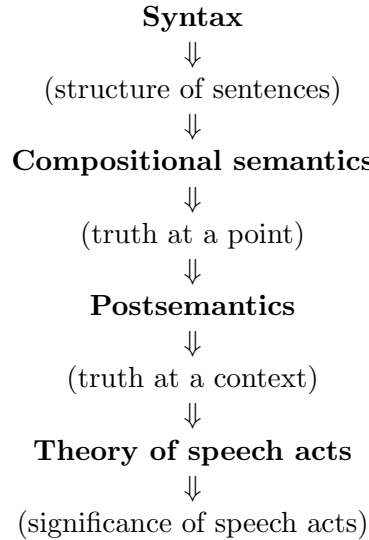
substance hangs on them. But the approach advocated here makes possible a kind of modularity that will be useful and illuminating later, when we introduce the third grade of truth relativity.

4 Compositional semantics and postsemantics

One way to get clearer about the distinction between truth at a point and truth at a context is to consider the different *roles* they play within a theory of meaning for a language. I find it helpful to divide such a theory into four parts: syntax, compositional semantics, postsemantics, and the theory of speech acts (see Figure 1). Syntax provides an account of the structure of sentences. Compositional semantics uses this as the basis for a recursive definition of truth at a point for declarative sentences. Postsemantics takes this definition as input and yields, as output, a definition of truth at a context of use for declarative sentences. This definition, in turn, is used by the theory of speech acts to give an account of the significance of the various speech acts that can be made through the *literal* use of declarative sentences or systematic transformations of them. (These transformations can include changes in mood and word order, deletion of words, substitution of question words like “who” or “what,” and insertion of phrases with a non-truth-conditional function, like “frankly” or “finally.” Thus, “Who is closing the window?” “Close the window,” and “Frankly, John is closing the window” all count as transformations of the declarative core “John is closing the window.”) A full theory of the use of language in communication would also need to account for nonliteral and extraliteral uses of language. The study of these falls to *pragmatics*, which I do not include as a part of the theory of meaning for a language, because its principles are not specific to any particular language.

The role of postsemantics in this scheme is to mediate between the compositional semantics and the theory of speech acts. The theory of speech acts will tell us, for example, that to make an assertion is to commit oneself to the truth of the asserted sentence at its context of use. So what we need to know in order to determine the significance of asserting a particular sentence is what a context of use must be like in order for the sentence to be true at that context. It won’t help to know at what *points* the sentence is true, unless we have some way of moving from truth at a point to truth at a context. Thus, the appropriate input for the theory of speech acts is a definition of truth at a context. However, truth at a context cannot be defined directly in the recursive clauses of a truth definition, except in the

Figure 1: Parts of a theory of meaning.



simplest languages. Once we have quantifiers, we will need to relativize truth to assignments, and assignments are not features of context. Similarly, if we have modal and temporal operators, we will have to relativize truth to world and time indices that can shift beyond the world and time of the context of use. If we are to have a compositional semantics, then, we must relativize truth to *points*. Postsemantics is needed to take us from this point-relative truth, which is not directly digestible by the theory of speech acts, to context-relative truth, which is.

Note that the output of the postsemantics is not a *function* from contexts of use to truth values, in the austere mathematical sense (a set of ordered pairs). The sentences “ $2 + 2 = 4$ ” and “There are infinitely many primes” would be associated with the same such function, since both are true at every possible context of use. Rather, the output of the postsemantics is a *definition* of truth at a context, or, in the case of a particular sentence S , a *specification* of the contexts at which S is true. The specification of the contexts at which “ $2 + 2 = 4$ ” is true can differ from the specification of the contexts at which “There are infinitely many primes” is true, even if these two sets of contexts turn out to be the same. The former might be “those contexts of use at which $2 + 2 = 4$,” while the latter might be “those contexts of use at which there are infinitely many prime numbers.” Because these specifications are distinct, the theory of speech acts can distinguish

between the assertoric significances of the two sentences.

In addition to truth at a context, the postsemantics may define falsity at a context, as well as validity and related logical notions. Whether a definition of falsity at a context is needed depends on the details of the theory of speech acts. On some accounts, the distinction between truth and untruth at a context is all that is needed; on others, falsity must be distinguished from other kinds of untruth.¹⁸ The arguments that follow will not depend on a resolution of this question.

Validity, in Kaplan's sense—preservation of truth at every possible context of use—is also a postsemantic idea.¹⁹ So are consistency, logical truth, and logical consequence, provided they are defined in terms of truth at a context of use. I do not deny that for some purposes—for example, in considering rules for free-variable reasoning or reasoning under a hypothesis—it may be useful to define validity via quantification over *points of evaluation* rather than contexts. And logical equivalence must certainly be defined via quantification over points, if logically equivalent formulas are to be intersubstitutable *salva veritate* in embedded contexts. However, the (postsemantic) notion of validity defined above precisely captures the idea that the valid inferences are the inferences we can *count on* to preserve truth, no matter how things are when we use them.²⁰

¹⁸For the former view, see Dummett's early article "Truth," reprinted in his *Truth and Other Enigmas* (Cambridge: Harvard University Press, 1978), 1–24. For the latter, see Ian Rumfitt, "'Yes' and 'No'," *Mind* 109 (2000), 781–823, at 796–800.

¹⁹See "Demonstratives," 547 (although Kaplan defines validity for sentences only, not for inferences), and "Afterthoughts," 594–7. Note that this characterization of validity remains schematic until we say exactly what is to count as a "possible context of use." I see no reason to think that there is only one correct or useful way to do this. In general, then, there will be *many* compatible notions of validity, each quantifying over a different set of contexts. There may also be alternative notions of validity that appeal to the preservation of some feature other than (or in addition to) *truth*: e.g. nonfalsity, warrant, or subject matter.

²⁰There is no need to quantify over *interpretations* in addition to contexts, because I am considering only fully interpreted languages (see note 2, above). In my view, it is a category mistake to think that a definition of validity for *inferences* should quantify over interpretations. Inferences consist of sentences, which already *have* interpretations. It is only when we are defining validity for inference *schemata* that quantification over interpretations becomes necessary, since schematic terms (the "nonlogical constants" of standard model theory) have no interpretation of their own. In languages without context sensitivity, the quantification over contexts is redundant, and validity (for inferences) reduces to simple material truth preservation. An inference in such a language is valid just in case it has a false premise or a true conclusion. Thus, the inference from " $2+2=5$ " to "there are only finitely many primes" is valid. It is a mistake to object that in order for the inference to be valid, there must be some *formal* relation between premises and

An example

Let us make this abstract framework a bit more concrete. Consider a language L_1 with quantifiers, truth-functional operators, tense operators (“*Was:*”, “*Will:*”), and the operator “*Now:*”. In order to accommodate the quantifiers, our points will need an assignment parameter and a domain parameter. In order to accommodate the tense operators, we will also need a *time* parameter. And in order to accommodate “*Now:*”, we will need a *time-of-use* parameter that can be held fixed while the time parameter varies.²¹ Truth at a point can then be defined as follows:

Compositional semantics for L_1 :

- An atomic sentence ‘ Fm ’ is true at $\langle D, a, t, t_u \rangle$ iff the extension of F at t contains the referent of m .²²
- An atomic open formula ‘ Fx_i ’ is true at $\langle D, a, t, t_u \rangle$ iff the extension of F at t contains the object assigned by a to x_i .
- ‘ $\neg\Phi$ ’ is true at $\langle D, a, t, t_u \rangle$ iff Φ is not true at $\langle D, a, t, t_u \rangle$.
- ‘ $\Phi \wedge \Psi$ ’ is true at $\langle D, a, t, t_u \rangle$ iff Φ is true at $\langle D, a, t, t_u \rangle$ and Ψ is true at $\langle D, a, t, t_u \rangle$.
- ‘ $\forall x_i\Phi$ ’ is true at $\langle D, a, t, t_u \rangle$ iff for every a' that differs from a at most in assigning a different object from D to x_i , Φ is true at $\langle D, a', t, t_u \rangle$.
- ‘*Was:* Φ ’ is true at $\langle D, a, t, t_u \rangle$ iff for some t' earlier than t , Φ is true at $\langle D, a, t', t_u \rangle$.
- ‘*Will:* Φ ’ is true at $\langle D, a, t, t_u \rangle$ iff for some t' later than t , Φ is true at $\langle D, a, t', t_u \rangle$.
- ‘*Now:* Φ ’ is true at $\langle D, a, t, t_u \rangle$ iff Φ is true at $\langle D, a, t_u, t_u \rangle$.

The postsemantics is as follows:

conclusion. The fact that an inference has a valid form—that is, is an instance of a valid inference *schema*—simply makes its validity more easily *knowable* (cf. Bertrand Russell, *Introduction to Mathematical Philosophy*, second edition [New York: Macmillan, 1920], 152–4). Generally speaking, our useful logical knowledge consists in knowledge of the validity of inference *schemata*, rather than particular inferences.

²¹The need for this kind of “double indexing” was first noticed by Hans Kamp, “Formal Properties of ‘Now,’” *Theoria* 35 (1971), 227–73.

²²I assume here that individual constants are temporally rigid, so we need not relativize their reference to times of evaluation. Predicates, on the other hand, are not temporally rigid, so we must relativize their extensions to times.

Postsemantics for L_1 : A sentence Φ is true [false] at a context of use C iff Φ is true [false] at every point $\langle D, a, t, t_u \rangle$ such that

- D = the domain intended by the speaker at C
- a = an assignment of values from the domain to the variables
- t = the time of C
- t_u = the time of C .

An inference is valid iff there is no context of use C such that the premises are all true at C but the conclusion is not true at C .²³

It is important to notice the different ways in which different parameters of points are treated in the postsemantics. Certain parameters (the domain, time, and time-of-use) have values *determined* by C , while others (the assignment) have values that are *constrained*, but not *determined* by C . (The assignment parameter is constrained by C , because it must assign objects from the contextually determined *domain* to the variables; but it is not determined by C , because there are many assignments that meet this constraint for any given context.) We have seen this difference before, when we noticed that the world parameter gets “initialized” by the context of use, while the assignment parameter must be quantified over. This difference in the way the postsemantics treats different parameters will play an important role in what follows.

5 Truth relative to a context of assessment

Up to this point, we have assumed that whether an assertion counts as literally true depends only on the sentence asserted (or some suitable expansion) and facts about the context of use (including relevant facts about the utterer’s intentions and the world in which the utterance occurs). This is a deeply embedded assumption in analytic philosophy of language: although truth for sentences must be relativized to a context of use, truth for utterances is supposed to be absolute.²⁴ Naive undergraduates may describe

²³Note that the quantification over contexts here brings with it the quantification over domains that is standard in definitions of validity for quantificational languages.

²⁴“Utterance” is ambiguous between the act of *uttering* and the object thereby produced; I use it in the “act” sense throughout. Thus a “sentence token”—a written in-

an assertion as “true for Joe but not for Sally,” but this is nonsense, or at best a misleading way of saying that Joe and Sally disagree.

My aim in this paper is to challenge this orthodox view. In the next section, I will present what seems to me a clear counterexample to the absoluteness of utterance truth. My aim in this section is more modest: I want to describe what it would be *like* to speak a language for which utterance truth is relative. I will first describe such a language in semantic terms. The truth of sentences in this language can depend not only on features of the context of use, but also on features of the context in which an utterance of the sentence is being *assessed*. Thus two kinds of linguistic context sensitivity are possible. A sentence is *use-sensitive* just in case its truth value varies with the context of use, keeping the context of assessment fixed, and *assessment-sensitive* just in case its truth value varies with the context of assessment, keeping the context of use fixed.

Having described the language semantically, I will consider what it would be like to *speak* such a language. In order to do this, I will need to provide an account of assertion that gives a role to truth relative to a context of assessment. I will then use this account to show that there is nothing *incoherent* about assessment sensitivity. Whether a particular language exhibits assessment sensitivity is a matter of contingent fact, to be settled by broadly empirical means.

Doubly relativized truth

Let us start, then, with L_1 from the previous section: the language with tense operators “*Will:*” and “*Was:*” and an operator “*Now:*” that forces its operand to be evaluated at the time of use. And let us add a new operator, “*Noy:*”, that forces its operand to be evaluated at the time of *assessment*. Call the resulting language L_2 . In order to state a recursive clause for “*Noy:*”, we will need a new parameter in our points: the time of assessment. The recursive clause then goes as follows:

Compositional semantics for *Noy:*

‘*Noy:* Φ ’ is true at $\langle D, a, t, t_u, t_a \rangle$ iff Φ is true at $\langle D, a, t_a, t_u, t_a \rangle$.

scription or acoustic wave pattern in the air—is not an utterance in my sense, and the fact that a sign on an office door saying “Back in an hour” can be accurate at one time and inaccurate at another is not a counterexample to the absoluteness of utterance truth (*pace* Philip Percival, “Absolute Truth,” *Proceedings of the Aristotelian Society* 94 [1994], 189–213, at 205). An utterance is a speech act, not an inscription or acoustic wave.

The interesting changes come in the postsemantics. Sentence truth is now relative to both a context of use and a context of assessment. Whereas the time and time-of-utterance parameters get initialized (as before) by the context of use, the time-of-assessment parameter gets initialized by the context of assessment.

Postsemantics for L_2 : A sentence Φ is true [false] at a context of use U and context of assessment A iff Φ is true [false] at every point $\langle D, a, t, t_u, t_a \rangle$ such that

- D = the domain intended by the speaker at U
- a = an assignment of values from the domain to the variables
- t = the time of U
- t_u = the time of U
- t_a = the time of A .

In addition, the definition of validity must be revised to take account of relativity of truth to contexts of assessment. There are at least two interesting notions of validity, one relative, the other absolute:

An inference is *valid relative to a context of assessment* A iff there is no context of use U such that the premises are all true at U and A but the conclusion is not true at U and A .

An inference is *absolutely valid* iff it is valid relative to every context of assessment.²⁵

It is vital not to confuse the relativization of truth to a context of assessment with relativization to a point of evaluation. All of the semantic examples we have considered so far have involved relativization of truth to points of evaluation, but only L_2 involves relativization of truth to a context of assessment. Indeed, contexts of assessment resemble contexts of use much more than they resemble points of evaluation. Like contexts of use, they are concrete contexts, not sequences of values for parameters. And like contexts of use, they play a role in the postsemantics, not the compositional semantics. The main difference is that contexts of assessment are associated with the *assessment*, rather than the *production*, of speech acts. It should be clear, then, that the relativization of truth to a context of assessment is

²⁵From here on, I will omit the definitions of validity in the postsemantics. The same definitions will work in all cases.

a *third* grade of truth relativity that goes beyond relativization to points of evaluation or to contexts of use.²⁶

Assertion and assessment sensitivity

The truth predicate defined by our postsemantics is now doubly relativized, to a context of use and a context of assessment. Since the output of the postsemantics is supposed to be the input for a theory of speech acts, we must show that this doubly contextual truth predicate can be put to use in an account of the literal significance of speech acts. But it is not at all obvious that it is suitable for this use. According to one standard account of assertion, the illocutionary point of an assertion is to commit the asserter to the truth of what is asserted.²⁷ In the idiom of sentences and contexts of use:

(Assertion-1) To assert a sentence *S* is (*inter alia*) to commit oneself to the truth of *S* (relative to the context of use).²⁸

But it is at least not obvious how one can commit oneself to the truth of a sentence that does not *have* an absolute truth value (even relative to a context of use). So it is not obvious that we can really understand what it would be to *assert* an assessment-sensitive sentence like “I am hungry noy.”²⁹

I think that we *can* understand what it is to assert an assessment-sensitive sentence, and without giving up the basic insight expressed by (Assertion-1). What is needed is a bit of reflection on what it *means* to commit oneself to the truth of a sentence (relative to a context of use).

²⁶The notion of context of assessment should also not be confused with Stefano Predelli’s notion of *context of interpretation* (“I am not here now,” *Analysis* 58 [1998], 107–115). Predelli argues that in making a recorded utterance, e.g. “I am not here now” on an answering machine, the speaker may *have in mind* a context of interpretation relative to which the context-sensitive expressions (“here,” “now”) are to be evaluated. My notion of context of assessment, by contrast, has nothing to do with the speaker’s intentions, and is not fixed in any way (even “intentionally”) by the context of use.

²⁷John Searle, *Expression and Meaning* (Cambridge: Cambridge University Press, 1979), 12.

²⁸The “*inter alia*” leaves room for the possibility that asserting a sentence involves more than simply committing oneself to its truth. Plausibly, the commitment must be undertaken publicly, by means of an overt utterance; perhaps there are other conditions as well. I focus here on commitment to truth, because it is the aspect of assertion that threatens to make assertion of assessment-sensitive sentences unintelligible.

²⁹For a discussion of a related criticism due to Gareth Evans, see section VI of my “Future Contingents and Relative Truth,” *Philosophical Quarterly* 53 (2003).

Some philosophers seem happy to take “commitment to truth” as basic. But I agree with Belnap *et al.*³⁰ that deontic constructions must take agentive complements, at the level of logical form. We can make sense of “being committed to Al Gore,” but only as meaning something like “being committed to *working for* (or perhaps *supporting*) Al Gore.” Similarly, when we say that Sarah is entitled to a cookie, we might mean that she is entitled to *take* a cookie, or perhaps to *eat* a cookie. When we say that Susy is committed to a house being built on this spot, we might mean that she is committed to building a house on this spot, or to seeing to it that others build one. But when no obvious agentive complement presents itself, we can’t make any sense of deontic constructions at all: what would it mean, for example, to be committed to the color of the sky?

So, in committing oneself to the truth of a sentence (at a context of use), what exactly is one committed to *doing*? Well, suppose I assert “Jake is in Boston.” If you ask “How do you know?” or challenge my claim more directly, by giving reasons for thinking it false—for example, “That can’t be: he has an engagement this afternoon in Los Angeles”—then it seems to me that I have an *obligation* to respond, by giving adequate reasons for thinking that my claim was true, or perhaps by deferring to the person who told me. If I can’t discharge this obligation in a way that meets the challenge, I must “uncommit myself” by retracting my assertion. If I neither withdraw the assertion nor reply to the challenge, I am shirking an obligation I incurred not *qua* moral agent or friend or member of polite society, but simply *qua* asserter.

These observations suggest an answer to our question. What I have committed myself to doing, in asserting that Jake is in Boston, is vindicating my claim if it is challenged.³¹ There may be no specific sanction for failing to follow through on this commitment. But if I fail too blatantly or too frequently, others may stop treating me as a being that is capable of undertaking this kind of commitment. They may still take my utterances as expressions of my beliefs, as we take a dog’s excited tail wagging as an ex-

³⁰ *Facing the Future*, 13.

³¹ For this way of looking at the commitment undertaken in asserting as a “conditional task responsibility” to vindicate a claim when it is challenged, see Robert Brandom, *Making it Explicit* (Cambridge: Harvard University Press, 1994), chapter 3, and “Asserting,” *Nous* 17 (1983), 637–50. There are some significant differences between the account proposed here and Brandom’s. Most prominently, Brandom does not use the notion of truth in his account; he is committed to explaining assertion independently of truth. Although I do not develop the idea in quite the same way as Brandom, I am much indebted to his work.

pression of its psychological state.³² They may even regard my utterances, if found to be reliable, as useful bits of information. But they will be treating me as a measuring instrument, not as an asserter. They will not take me to be *committing myself* to the truth of my utterances.

If this is the right way to think of “commitment to the truth of a sentence,” then we may replace (Assertion-1) with the following:

(Assertion-2) To assert a sentence S (at a context U) is (*inter alia*) to commit oneself to providing adequate grounds for the truth of S (relative to U), in response to any appropriate challenge, or (when appropriate) to defer this responsibility to another asserter on whose testimony one is relying. One can escape this commitment only by withdrawing the assertion.

A few points of clarification are in order.

1. Not every challenge counts as “appropriate.” A mathematics professor who has asserted that it is impossible to trisect an angle using a compass and straightedge need not respond to every crackpot counterexample she receives in her e-mail. Whether a challenge counts as “appropriate” will depend on what was asserted, who asserted it, who is challenging it, what grounds for challenge have been offered, and the dialectical context. I will not attempt to work out criteria for appropriateness here, but it seems to me that we do distinguish appropriate from inappropriate challenges in our practice.
2. In speaking of “grounds for the truth of S (relative to U),” I do not want to imply that the grounds must be explicitly metalinguistic and semantic. “It had gills” can count as grounds for the truth of the sentence “It was a fish.” In general, an utterance of sentence S_1 at context C_1 counts as grounds for the truth of a sentence S_2 relative to C_2 iff (a) S_1 is true at C_1 and (b) the truth of S_1 at C_1 entails, or is evidence for, the truth of S_2 at C_2 .
3. What grounds count as “adequate” will depend on a number of factors, including the nature of the challenge and what the challenger knows. In response to some challenges, weak inductive evidence might be adequate; for others, demonstrative proof might be required. In addition,

³²Note that this is a less exacting notion of “expressing a belief” than the one used by many speech-act theorists: see Kent Bach and Robert M. Harnish, *Linguistic Communication and Speech Acts* (Cambridge: MIT Press, 1979), 15.

grounds are adequate only if the challenger is in a position to recognize that they *are* grounds, or would be in a position to recognize this if she had an explicit understanding of the semantics of the relevant sentences. Thus, an utterance of “The left and right footprints are of different sizes” can count as grounds for the truth of “The butler did it” if the challenger knows that the butler’s feet are of different sizes, but not otherwise.

4. I leave it open when (if ever) it is appropriate to defer the responsibility for responding to a challenge to another asserter. This question is connected with delicate issues about the epistemology of testimony. Certainly it is not appropriate to defer to a speaker whom one has good reason to think is untrustworthy, or whose claim one has good reason to think is false. But when these defeating conditions are absent, it may be appropriate to defer instead of responding directly to a challenge.³³
5. Withdrawal may be formal (“I take that back”) or informal (“So I was wrong”). In some cases it may not need to be made explicit at all, as there is an expectation that an assertion proven to be false will be withdrawn. In those cases, one must speak up if one wishes to remain committed to the claim, or one will be presumed to have withdrawn it.

Note that 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, as long as that assertion remains unchallenged. Thus, on this account, one can lie without violating the constitutive norms of assertion (of course, one may be violating other, moral and prudential norms).³⁴

Once we have cashed out (Assertion-1) as (Assertion-2), it is easy to find a role for truth relative to a context of assessment. A slight modification will allow us to make sense of the assertion of assessment-sensitive sentences:

³³See Brandom, *Making It Explicit*, 174–7.

³⁴Contrast the “knowledge account of assertion” defended by Timothy Williamson (*Knowledge and Its Limits* [Oxford: Oxford University Press, 2000], ch. 11) and others. Of course I owe an explanation of the phenomena the knowledge account of assertion purports to explain, primarily the infelicity of asserting “*P*, but I don’t know that *P*.” But I have the resources for such an explanation. In asserting this sentence, one is committing oneself to justifying *P* if it is challenged, while conceding that one cannot meet certain legitimate challenges. Indeed, one is committing oneself to *showing* that one cannot meet certain legitimate challenges to *P*. It is certainly infelicitous to commit oneself simultaneously both to performing a task and to showing that one cannot perform this task.

(Assertion-3) To assert a sentence S (at a context U) is (*inter alia*) to commit oneself to providing adequate grounds for the truth of S (relative to U and one’s current context of assessment) in response to any appropriate challenge, or (when appropriate) to defer this responsibility to another asserter on whose testimony one is relying.³⁵ One can escape this commitment only by withdrawing the assertion.

On this account, when an assertion is challenged, the asserter must respond by providing adequate grounds for the truth of her utterance relative to her *current* context of assessment—the context she is in at the time of the challenge—even if this is very different from the context in which the assertion was made.³⁶

Note that this account presupposes only that it makes sense to talk about contexts of assessment; it does not presuppose that there *is* any assessment sensitivity. Thus it ought to be acceptable even to an opponent of assessment sensitivity, who will not make use of the context of assessment in the postsemantics, and who will thus regard “true at U and one’s current context of assessment” as a long-winded way of saying “true at U .” I do not see how one could object to the very notion of a context of assessment. Assessing speech acts is certainly something we do: how could it not make sense to talk of the context in which we do this on a particular occasion?

Using (Assertion-3), we can see what it would be like to assert an assessment-sensitive sentence. Suppose I assert “Jane is hungry now” at t_1 , when Jane is in fact hungry. Relative to context of use t_1 and context of assessment t_1 , my assertion is manifestly true.³⁷ If it is challenged, I can easily justify it. But then at t_2 Jane eats a big meal and becomes stuffed. Relative to the new context of assessment (t_2) and the context of use (still t_1), my assertion is manifestly untrue, so (upon being challenged) I am obliged to withdraw it. Of course, if I do not withdraw my assertion, this

³⁵When I defer this responsibility to someone else, does she have to give grounds for the truth of S relative to U and *my* current context of assessment, or relative to U and *her* current context of assessment? Presumably the latter, but I am not sure what to say here.

³⁶Again, in speaking of “grounds for truth” I do not mean to imply that the justification must be explicitly semantic. An utterance of sentence S_1 at context of use C_1 counts as grounds for the truth of a sentence S_2 relative to context of use C_2 and context of assessment A if the truth of S_1 relative to C_1 and A entails, or is evidence for, the truth of S_2 relative to C_2 and A .

³⁷Since the only relevant features of contexts here are times, I simplify by identifying the context of use with the time at which the assertion is made and the context of assessment with the time at which the assertion is being assessed.

obligation will lapse as soon as Jane becomes hungry again (at t_3), because I will then be able to show that the assertion is true, relative to its context of use (t_1) and my current context of assessment (t_3).

Admittedly, it is hard to see what the *point* of asserting a sentence like “I am hungry noy” would be. It is not surprising that we do not have a word that works like “noy” in English. But it is enough for now that we understand what it would *be* to assert such a sentence. Once we have cashed out the idea of “commitment to truth” in properly agentive terms, we can make good sense of commitment to the truth of an assessment-sensitive sentence. To be committed to the truth of a sentence asserted at context U is simply to be committed to producing adequate grounds for its truth (relative to U and one’s current context of assessment) when appropriately challenged (or to defer this responsibility, if appropriate). If this obligation cannot be met—for example, if the sentence is shown decisively to be untrue (again, relative to U and one’s current context of assessment)—then one must withdraw the assertion. This commitment is no more complex or hard to understand than a commitment to empty the trash can whenever it is full (or, if one cannot do this, to call a janitor), or a commitment to eat lunch every day at noon in whatever time zone one happens to be in that day.

It appears, then, that a postsemantics that defines doubly relativized truth *can* serve as input to an independently plausible account of assertion. It has not been shown, however, that any sentence of English (or any other real language) exhibits assessment sensitivity.³⁸ Asserting “I am noy hungry” may be coherent, but it seems pointless. In the next section, however, I will display what seems to me a convincing example of assessment sensitivity in English: discourse about future contingents.

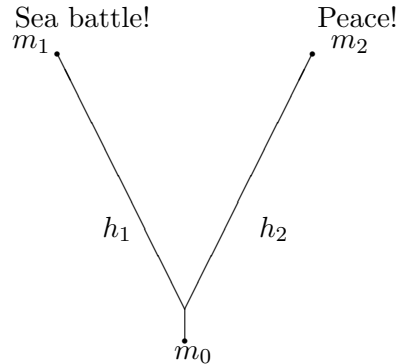
6 Future contingents

Suppose that the future is genuinely undetermined. As of now, there is a real possibility that there will be a sea battle tomorrow and a real possibility that there will not be one. It’s not just that I don’t know which will happen, or that I can’t know, or even that I *couldn’t* know even with idealized information about the present state of the world, but that as a matter of

³⁸It has been suggested to me that we can get the effect of saying “I am hungry noy” in English by saying something like “Always consider me to be asserting that I’m hungry.” But that is an imperative and could not have the same illocutionary significance as an assertion. Moreover, the imperative instructs one to take the speaker to be making a *series* of assertions, at different times, of “I am hungry.” That is different from asserting, at one time, “I am hungry noy.”

objective fact, either might happen.³⁹ We can model the situation with a picture (Figure 2). In this picture, m_0 , m_1 , and m_2 are “moments”: think

Figure 2: Branching histories.



of them as concrete time-slices of the universe. h_1 and h_2 are “histories,” which you can think of as maximal chains of moments, partially ordered in a “tree” structure by a causal/temporal relation $>$.⁴⁰ There can, in general, be multiple histories running through a single moment, as h_1 and h_2 both run through m_0 .⁴¹ I’ll assume that m_1 and m_2 happen one day after m_0 . At m_1 , there is a sea battle, while at m_2 , there is not one.

Now suppose that a speaker at m_0 asserts “There will be a sea battle

³⁹This notion of indeterminism is stronger than what is sometimes called “causal indeterminism.” The future is causally indeterministic just in case it could not be predicted even with ideal knowledge of the present and past state of the universe and all laws of nature. Objective indeterminism implies causal indeterminism, but not vice versa: one could, in principle, have an objectively deterministic world that is not lawlike enough to be causally deterministic (for example, a world with pre-ordained miraculous interventions).

⁴⁰The partial order is a *tree* because we assume “no backward branching”: every moment has a determinate past. For more details on the semantico-metaphysical picture I’m appealing to here, see Belnap et al., *Facing the Future*, 29–32, 139–141. For a relativistically acceptable version of branching histories, see Nuel Belnap, “Branching Space-time,” *Synthese* 92 (1992), 385–434. For reasons for preferring the branching histories framework to the “times cross worlds” ($T \times W$) framework that is more typically used for combining tense and modality, see *Facing the Future*, 196–209, 231–3. For arguments against branching, see David Lewis, *On the Plurality of Worlds* (Oxford: Basil Blackwell, 1986), 206–9. Lewis’s arguments are effectively countered by Belnap et al. in *Facing the Future*, 207–9.

⁴¹Strictly speaking, I should say: a single moment can *be a member of* many histories. But I’ll continue to use the more intuitive terminology of histories “running through” moments.

tomorrow.” How should we assign a truth value to that assertion? This is the age-old problem of future contingents. It is age-old because none of the answers on the traditional menu—true, false, or neither—seem completely satisfactory. In this section, I will argue that future contingents are in fact assessment-sensitive, and that is why they have stubbornly resisted treatment in orthodox semantic frameworks.⁴² (Note that it is not a solution to the semantic problem to deny that the world *is* objectively indeterminate. Perhaps it is objectively determinate, but that is not for the semanticist to decide. Here I do not presuppose that the world is objectively indeterminate, but only that our semantic theories ought to leave *room* for this possibility. The theory of meaning should not presuppose either that the world is indeterministic or that it isn’t.)

Let’s start with the compositional semantics. By taking points of evaluation to contain both a moment and a history parameter, we can give a simple and attractive recursive semantics for tense operators and historical modalities.⁴³

**Compositional semantics for “it will be the case that”
and “it is settled that”:**

- ‘*Will*: Φ ’ is true at $\langle m/h \rangle$ iff for some $m' > m$ on h , Φ is true at $\langle m'/h \rangle$.
- ‘*Sett*: Φ ’ is true at $\langle m/h \rangle$ iff for every h' through m , Φ is true at $\langle m/h' \rangle$.

These operators behave just as you would expect they would: ‘*Will*:’ shifts you forward along a history, while ‘*Sett*:’ quantifies over all histories through a moment.

I now want to show that unless we allow assessment sensitivity, we cannot complete this compositional semantics with a reasonable postsemantics. Suppose that there is no assessment sensitivity in our discourse about future contingents. Then our postsemantics must define *truth at a context of use* in terms of *truth at a moment/history pair*. Let us see how this can be done. The moment parameter can be initialized by the context of use (it is what I called a *determined* parameter in section 4, above). But what about the history parameter? It is not determined, but *constrained*. Suppose that at

⁴²For a fuller discussion, see “Future Contingents and Relative Truth,” 321–36.

⁴³It is assumed that in every point $\langle m, h \rangle$, $m \in h$ (that is, h “passes through” m). Following Belnap, I use the notation m/h to bring out this connection. For simplicity of exposition, I suppress other parameters, e.g., domain, assignment, speaker, moment of use, interpretation, and structure.

m_0 I assert “There will be a sea battle.” There are *two* histories (h_1 and h_2) going through the moment of use, and nothing about the objective situation or the speaker’s intentions singles out one of them. There is a “moment of the context of use,” but not a “history of the context of use.”⁴⁴

The thin red line

It is tempting here to suppose that the context of use *does* in fact determine a unique privileged history—an “actual future,” which we might think of as marked out by a “thin red line” on our tree—but I agree with Belnap and Green that this amounts to giving up on objective indeterminism.⁴⁵ On this picture, the non-red branches in the tree are supposed to represent objectively possible futures. But their non-redness indicates precisely that they *won’t* be the continuations of the history that includes the utterance in question. Looking down on the tree of histories, God could see that given the past and the context of use, only one continuation remains in play: the one marked with the thin red line. In what sense, then, are the others “possibilities”? They are possible in an *epistemic* sense: the utterer doesn’t know which history is marked out with the thin red line. But objectively speaking, they are not genuine possibilities.

Of course, we *do* often say that one of two objectively possible outcomes turned out to have been the actual one, and even that it “was going to be” the actual one. “At the time, the coin could have landed either Heads or Tails, but as we now know, it was actually going to land Heads.” It is this kind of talk that makes the thin red line seem intelligible and even compelling. But we can make this kind of claim only from a particular perspective *within* the tree of branching histories. For example, speakers at m_1 (in Figure 2) can truly say that a sea battle “was actually going to happen” at m_0 , while speakers at m_2 can truly say that a sea battle “was not actually going to happen.”⁴⁶ *Qua* semanticists, however, we do not speak from the perspective of any particular moment on the tree; instead, we take a God’s eye point of view, looking down on the tree from the outside, and

⁴⁴See *Facing the Future*, 151–2.

⁴⁵Cf. Nuel Belnap and Mitchell Green, “Indeterminism and the Thin Red Line,” *Philosophical Perspectives* 8 (1994), 365–88, revised as Belnap et al., *Facing the Future*, ch. 6. Belnap and Green also give some purely *technical* arguments against the use of a thin red line in the compositional semantics for tense operators, but these arguments have no force against an invocation of a thin red line in the *postsemantics*, of the sort I am envisioning here.

⁴⁶See below for a discussion of the semantics of “actually” in an indeterministic framework.

try to say how the truth of sentences depends on features of the context of use described without using indexicals.⁴⁷ From the “external” semanticist’s perspective, 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 the internal and external perspectives simultaneously, that it can seem to make sense to mark out one of the histories with a thin red line.

Supervaluation

So it seems we have no choice but to *quantify* over the history parameter (just as we had to quantify over the assignment parameter in L_1). The result is the following *supervaluational* post-semantics:

Supervaluational post-semantics for branching histories:

Φ is true [false] at a context of use C iff Φ is true [false] at every point $\langle m/h \rangle$ such that

- m = the moment of C
- h runs through the moment of C .

On this view, a sentence is true at a context iff it is true on all histories that are open possibilities at the context of use, false iff it is false on all such histories.⁴⁸

It is not hard to see what is wrong with the supervaluational approach. Consider again Figure 2. Suppose that at m_0 Jake asserts “There will be a sea battle within 24 hours.” According to supervaluational postsemantics, Jake’s assertion is neither true nor false, since there is one history through m_0 on which there is a sea battle the next day (h_1) and one on which there is not (h_2). But now put yourself in the shoes of someone at m_1 . A sea battle is raging around you. Wouldn’t you say that Jake’s assertion turned out to be *true*? After all, he asserted that there would be a sea battle, and there was one. Yet supervaluational postsemantics says that his assertion is *not* true.⁴⁹

⁴⁷This constraint rules out clauses like “*Will*: Φ ’ is true (if uttered now) iff Φ will be true,” as well as “The name ‘Berkeley’ refers to the city in which I am now located.”

⁴⁸See R. H. Thomason, “Indeterminist Time and Truth-value Gaps,” *Theoria* 36 (1970), 264–81.

⁴⁹Some care is needed here: the spectator at m_1 is employing an object-language truth predicate, not the semanticist’s metalinguistic truth predicate. Without some story about the relation between these, we cannot be sure that his truth evaluation is incompatible with the supervaluationist’s. But surely a plausible constraint on the semantics for an

Moreover, when combined with a minimal assumption about assertion—that one is obliged to withdraw assertions that have been shown to be untrue—supervaluational postsemantics predicts that Jake ought to withdraw his assertion if presented with compelling proof that it was not *settled* at the time of his assertion that there would be a sea battle. Thus, he can be obliged to withdraw his assertion, on grounds of its proven untruth, *even if there turns out to be a sea battle!* This result seems clearly wrong. There is a difference between asserting that *it is settled or determined* that there will be a sea battle tomorrow, and simply asserting that there will be a sea battle tomorrow. But supervaluationism seems to run these together. As Dummett points out, on a supervaluational semantics for future contingents, “. . . the force of uttering ‘*Sett:A*’ is no greater than that of asserting *A*, and the conditions for the two assertions to be correct will coincide.”⁵⁰ Relatedly, the supervaluational account makes the inference from *A* to ‘*Sett:A*’ come out valid. For *A* can be true at a context of use only if it is true at the moment of use and all histories passing through this moment, and at any such context of use ‘*Sett:A*’ is also true. This is, to say the least, counterintuitive.

But what alternative to supervaluational postsemantics is there? We have to do *something* with the history parameter, and if context does not supply a privileged history, what can we do but quantify?

Assessment sensitivity to the rescue

I suggest that the only way to get a correct postsemantics for future contingents in the branching histories framework is to relativize truth to a context of assessment as well as a context of use.⁵¹ Intuitively, Jake’s assertion is true as assessed at m_1 , false as assessed at m_2 , and neither true nor false as assessed at m_0 . And this is just the result we get from the following

object-language truth predicate (in a framework without assessment sensitivity, like the supervaluationist’s) is that an utterance u of S at C should fall into the extension of “true” just in case S is true at C . Given this constraint, there is indeed a conflict between the supervaluationist’s truth evaluation and the intuitive truth evaluation of a spectator at m_1 . Thanks to Lionel Shapiro for pressing me to clarify this point.

⁵⁰*The Logical Basis of Metaphysics* (Cambridge: Harvard University Press, 1991), 80. I have altered the quotation slightly, using “*Sett:A*” instead of Dummett’s “*UA*”. Note that Dummett is not saying that ‘*Sett:A*’ and *A* are *equivalent*: he concedes that they can differ in their effects on the truth conditions of sentences containing them.


⁵¹There is an anticipation of this view in Dummett: “We might consider the truth of a tensed sentence as relativized to two times, the time of utterance and the time of assessment” (*Frege: Philosophy of Language*, second edition [Cambridge: Harvard University Press, 1981], 395).

Alternative post-semantics for branching histories: Φ is true [false] at a context of use U and context of assessment A iff Φ is true [false] at every point $\langle m/h \rangle$ such that

- $m =$ the moment of U
- h passes through m and (if the moment of $A \geq m$) through the moment of A as well.

The essential structural feature of this account is what Nuel Belnap calls “double time references.”⁵² We evaluate Φ with respect to the moment of use and the histories passing through the moment of *assessment* (see Figure 3).⁵³

Figure 3: Double time references.



⁵²See *Facing the Future*, 175, and “Double Time References: Speech-act Reports as Modalities in an Indeterministic Setting” (*Advances in Modal Logic*, vol. 3, ed. F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyashev [Stanford: CSLI, 2001], 1–22). Note that Belnap does not relativize truth to a context of assessment. Realizing that there is no good way to assign “absolute” truth values to future contingents, he chooses to dispense entirely with what I have called “postsemantics” in his account of tensed sentences and appeals directly to point-relative truth in his account of assertion. The basic idea is that asserting Φ effects a change in normative status such that one is “vindicated” at every later moment at which Φ becomes settled true (with respect to the context of use), and “impugned” at every later moment at which Φ becomes settled false (with respect to the context of use). (Note the “double time reference” idea.) An advantage of the present proposal, which interposes the postsemantics between the compositional semantics and the account of assertion, is that it allows the account of assertion to deal with a uniform notion of truth at a context of utterance and context of assessment, thereby screening it off from the messy details of points and thus from the particular expressive resources of the language. Unlike Belnap’s account of assertion, (Assertion-3) makes no reference to moments, histories, or anything else that is specific to tensed language.

⁵³This assumes that the moment of assessment occurs *after* the moment of use in the historical ordering. If not, then we look at all of the histories passing through the moment of use.

The alternative postsemantics shows neatly how assertions of

- (1) There will be a sea battle tomorrow

and

- (2) It is settled that there will be a sea battle tomorrow

undertake different assertoric commitments. Whereas an assertion of (1) at m_0 is true as assessed from m_1 , an assertion of (2) at m_0 is false as assessed from m_1 . The reason is that the “It is settled that” operator instructs us to consider *all* of the histories passing through the moment of evaluation (here, initialized by the context of use as m_0). Since “there will be a sea battle tomorrow” is false at m_0/h_2 , “It is settled that there will be a sea battle tomorrow” is not true even at m_0/h_1 . So an assertion of (2) at m_0 comes out false whether the moment of assessment is m_0 , m_1 , or m_2 . The fact that there is a context of utterance and context of assessment at which (1) is true and (2) false shows that the inference from (1) to (2), which comes out valid on the supervaluational semantics, gets counted here (rightly) as invalid.

We can also explain the infelicity of asserting

- (3) There will be a sea battle tomorrow, but it is possible that there will not be a sea battle tomorrow,

where “possible” is an objective, not an epistemic modality, equivalent to “not settled that it is not the case.”⁵⁴ On the proposed semantics, this sentence is not contradictory; it is *true*, for example, as uttered at m_0 and assessed at m_1 . Nonetheless, it would be very odd to assert this sentence at m_0 , because if one’s assertion were challenged at m_0 , adequate grounds for the truth of either conjunct would establish the falsity of the other. For an adequate justification of the first conjunct would require showing that “There will be a sea battle tomorrow” is true at m_0 on *all* histories through the moment of assessment (here, m_0), whereas an adequate justification of the second would require showing that this same sentence is not true at m_0 on all histories through m_0 . Of course, if the assertion is not challenged at m_0 , it can stand, and at m_1 there will be no trouble meeting a challenge. Still, there is something odd about making an assertion that *manifestly* cannot be defended at the moment it is uttered, even if it might become defensible later.

⁵⁴For a discussion of the parallel case with epistemic possibility, see Keith DeRose, “Simple ‘Might’s, Indicative Possibilities and the Open Future,” *Philosophical Quarterly* 48 (1998), 67–82.

Semantics for “actually”

Admitting assessment sensitivity also makes possible an elegant semantics for “actually” in a branching histories setting. This is important not just for its own sake, but because it will allow us to block one of the most powerful motivating arguments for a “thin red line:” the fact that we readily say things like “Though it was possible that the coin would land Tails, it was actually going to land Heads.” The semantics I will offer shows how an utterance of this sentence can be true even without a thin red line.

In setting up the problem, it is useful to recall Lewis’s treatment of actuality. Lewis argues—correctly, in my view—that actuality is an indexical concept.⁵⁵ It’s not that out of all the possible worlds, one has the special property of being *actual*; rather, to say that a world is actual is to say that it is *our* world, the world of the context of use. In order to give a recursive definition of truth at a point for a language including an actuality operator, then, one needs an (unshiftable) “world of the context of use” parameter (w_u). Then the following simple definition will do:⁵⁶

Lewis semantics for “actually”: ‘*Actually*: Φ ’ is true at $\langle w_u, t_u, w, t \rangle$ iff Φ is true at $\langle w_u, t_u, w_u, t \rangle$.

In English: ‘*Actually*: Φ ’ is true at an arbitrary world just in case Φ is true at the world of the context of use.

Lewis’s semantics respects an important feature of actuality operators, which I call

Initial Redundancy: An operator \star is initial-redundant just in case for all sentences Φ , ‘ $\star\Phi$ ’ is true at exactly the same contexts of use and assessment as Φ (equivalently: the inference from each to the other is absolutely valid).

If “*Actually*:” were not initial-redundant, it might sometimes happen that one could truly assert a sentence S , but not ‘*Actually*: S ’ (or perhaps vice versa). But that does not seem to be possible. Whenever we assert that something is the case, we could just as well have asserted that that same thing is *actually* the case. (That is not to say that actuality operators are redundant *everywhere*; they can make a big difference in *embedded* contexts.) It is easy to verify that Lewis’s actuality operator is initial-redundant, provided we assume (as Lewis does) that the context of use determines a unique

⁵⁵ *On the Plurality of Worlds* (Oxford: Basil Blackwell, 1986), 92–6.

⁵⁶ I leave out the domain and assignment parameters in order to focus on the key issue.

world and time of use.⁵⁷ However, this is not an assumption we can accept in a branching histories framework. Lewisian worlds have unique histories, so the assumption that the context of use determines a unique world amounts to the assumption that it determines a unique future history.

How, then, can we adapt Lewis’s insight about the indexicality of “actually” to a branching histories framework? We cannot simply replace Lewis’s “world of the context of use” parameter with a “history of the context of use” parameter, for there is in general no unique history of the context of use. We can, however, introduce parameters for the moment of use (m_u) and the moment of assessment (m_a) (these will be initialized in the postsemantics by the context of use and the context of assessment, respectively). Given our postsemantics, Initial Redundancy imposes the following constraint on an actuality operator:

Constraint on the semantics for “actually”: ‘*Actually*: Φ ’ should be true at just those contexts of use and assessment at which Φ is true—that is, at just those contexts of use and assessment such that Φ is true at the moment of use on all histories through the moment of assessment (or, if the moment of assessment is not greater than the moment of use in the historical ordering, on all histories through the moment of use).

There are at least two operators that satisfy this constraint:⁵⁸

⁵⁷Let U be a context of use and w_u and t_u the world and time of U . Then ‘*Actually*: Φ ’ is true at U iff ‘*Actually*: Φ ’ is true at the point $\langle w_u, t_u, w_u, t_u \rangle$ iff Φ is true at the point $\langle w_u, t_u, w_u, t_u \rangle$ iff Φ is true at U .

⁵⁸It is interesting to compare this treatment of the actuality operator with Belnap’s, in *Facing the Future*, 246–7. Belnap defines two operators (adjusting for my somewhat different notation):

‘*Actually*₁: Φ ’ is true at $\langle m_u, m_a, m/h \rangle$ iff for every history h' running through m_u , Φ is true at $\langle m_u, m_a, m_u/h' \rangle$.

‘*Actually*₂: Φ ’ is true at $\langle m_u, m_a, m/h \rangle$ iff either (a) h runs through m_u and Φ is true at $\langle m_u, m_a, m_u/h \rangle$, or (b) h does not run through m_u and for every history h' running through m_u , Φ is true at $\langle m_u, m_a, m_u/h' \rangle$.

(Belnap does not use a moment-of-assessment parameter, of course, but I include it to ease comparison.) “*Actually*₁:” is not initial-redundant, at least on the postsemantics I have defended. For example, if it was not *settled* yesterday that there would be a sea battle today, but there turns out to be a sea battle today, then an assertion yesterday of “*Actually*₁:There will be a sea battle tomorrow” counts as false (assessed from today), while an assertion yesterday of “There will be a sea battle tomorrow” counts as true (assessed from today). “*Actually*₂:” avoids this problem, because the postsemantics pays attention *only* to points in which (a) the history parameter runs through the moment of

Compositional semantics for “actually” and “actually now”: ‘*Actually*: Φ ’ is true at $\langle m_u, m_a, m/h \rangle$ iff for every history h' running through m (and, if $m_a \geq m$, through m_a as well), Φ is true at $\langle m_u, m_a, m/h' \rangle$.

‘*Actually-Now*: Φ ’ is true at $\langle m_u, m_a, m/h \rangle$ iff for every history h' running through m_u (and, if $m_a \geq m_u$, through m_a as well), Φ is true at $\langle m_u, m_a, m_u/h' \rangle$.

The difference between ‘*Actually*:’ and ‘*Actually-Now*:’ is that the former looks at the truth values of the embedded sentence at the moment of *evaluation* and the histories through the moment of assessment, the latter at the moment of *use* and the histories through the moment of assessment.

Note that this semantics allows

- (4) It was possible that the coin would land Tails, though it was actually going to land Heads (*Was*:*[Poss*:*Will*:*Heads* \wedge *Actually*:*Will*:*Tails*])

to be true, relative to a context of use at which the coin has landed Tails (and any context of assessment). In denying that there is a “thin red line,” then, we need not deny that this sentence can be truly asserted.

Semantics for “True”

So far, we have studied assessment sensitivity from an “external” semanticist’s point of view, using the doubly relativized notion of *truth at a context of use and a context of assessment*. But it is worth asking how assessment-relative truth will look from an “internal” point of view, to the *participants* in a linguistic practice. Because they always occupy a particular context of assessment, the participants are not likely to refer explicitly to contexts of assessment in their appraisal of utterances as “true” or “false.” They are likely to say things like “Billy’s assertion is true,” using a one-place truth predicate. How should we understand such claims, if truth is assessment-relative? Relatedly, how will participants in a linguistic practice *express* the

use and (b) the moment parameter = the moment of use. On all such points ‘*Actually*₂: Φ ’ will have the same truth value as Φ , so the two will be true (or false) at exactly the same contexts of use and assessment. Initial redundancy is secured—but at the cost of a complex and seemingly *ad hoc* semantics. Like ‘*Actually-Now*:’, ‘*Actually*₂:’ looks at the moment of use, rather than the moment of evaluation. A salient difference between them is that the latter, but not the former, is sensitive to the history parameter. As a result, ‘*Poss*:*Actually*₂: $\Phi \wedge$ *Poss*:*Actually*₂: $\neg\Phi$ ’ can be true at a context of use and context of assessment, while ‘*Poss*:*Actually-Now*: $\Phi \wedge$ *Poss*:*Actually-Now*: $\neg\Phi$ ’ cannot be. In this respect “*Actually-Now*:” comes closer to Lewis’s actuality operator.

assessment sensitivity of an assertion? Will it be correct for them to say things like “Billy’s assertion wasn’t true before, but now it is true?” Will they be able to formulate the norm for withdrawal of assertions from an “internal” point of view, without talking explicitly of contexts of assessment?

To answer these questions, we need to give a semantics for an object-language truth predicate applicable to utterances. Or rather, we need to give a *naive* semantics for object-language “true.” The point here is to show how “true” should be handled in a language with assessment-sensitive expressions, not to deal with the semantic paradoxes. So I won’t worry here about the threat of inconsistency.

Intuitively, an utterance of ‘*u* is true’ ought to be true just in case the sentence *S* of which *u* is an utterance is true relative to the context at which it was uttered. This idea can be extended in a natural way to a framework with doubly-relativized truth:

Constraint on the semantics for “True” : An utterance *u* of sentence *S* at context of use C_1 should fall under the extension of “True” relative to context of use C_2 and context of assessment *A* just in case *S* is true relative to C_1 and *A*.⁵⁹

This constraint ensures that ‘*u* is true’ will be assessment-sensitive just in case *u* is an utterance of an assessment-sensitive sentence.⁶⁰ It also ensures that “true” can be used in an “internal” formulation of the norm for withdrawal of assertions: “Whenever your assertion is appropriately challenged, you must give adequate grounds for thinking that it is true, or withdraw it.” Anyone following this norm correctly will conform to (Assertion-3). For where *u* is an assertion of *S* by *A* at C_0 , adequate grounds for the truth of *u* (assessed at C_1) will be adequate grounds for the truth of *S* relative to context of use C_0 and context of assessment C_1 .

It should be clear that the constraint itself cannot serve as a compositional semantics for “True.” Truth relative to a context of use and context of assessment is defined in the postsemantics in terms of truth at a point; if we use it in our compositional semantics to define truth at a point, we make

⁵⁹Compare the constraint in note 49, above.

⁶⁰An alternative proposal would use the context of *use* of ‘*u* is true’ as the context of assessment for the sentence of which *u* is an utterance. On this proposal, “true” would be use-sensitive, not assessment-sensitive. However, I think the proposal in the main text comes closer to capturing the way we use “true” in English. It would sound very strange to assert: “*u* is not true, but I stand by what I said yesterday when I said, ‘*u* is true.’” But if “true” is use-sensitive, this assertion should be no stranger than “It is not cloudy now, but I stand by what I said yesterday when I said, ‘It is cloudy now.’”

our definitions circular. To get around this problem, we need to include parameters in our points that get initialized by the context of assessment. In the case of the simple tensed language we have been considering, all we need is a parameter for the moment of assessment, which we had already introduced to handle “actually.” We can then give the following semantics:

Compositional semantics for “True”:⁶¹ “True” is true of x at $\langle m_u, m_a, m/h \rangle$ iff for some sentence S and moment m' , x is an assertion of S at m' , and for all h' running through m' (and, if $m_a \geq m'$, through m_a as well), S is true at $\langle m', m_a, m'/h' \rangle$.

Note that the last part of the definition mirrors the definition of truth at a context of utterance and context of assessment in the postsemantics. In this way, we get a semantics for “True” that conforms to the constraint, but without courting circularity.

Because “True” is not sensitive to the moment of evaluation, it cannot be used to make explicit the relativity of the truth of future contingents to a context of assessment. It is never true to say of a single assertion—even an assertion of “There will be a sea battle tomorrow”—that “it is True now, but was not True before.” In this respect “True” seems to match our use of “true” in English. It would sound decidedly odd to say “Bill’s assertion is true now but it was not true when it was made.” Indeed, we normally don’t say that a (past) utterance “is” true at all; we let the time of the utterance determine the tense of the copula and say that it “was” true. If “true” were sensitive to the moment of evaluation, one would expect that both “is true” and “was true” would be used.

Perhaps the fact that we cannot truly say that an utterance of a future contingent was not True before, but is True now, explains the intuitive pull of the doctrine that utterance truth is absolute. We confuse invariance across moments of evaluation with invariance across moments of assessment. Carefully distinguishing between the three grades of truth relativity is thus a crucial first step in the defense of relativism about utterance truth.

7 Accommodation

In section 5, I argued that there was nothing incoherent about relativizing truth to a context of assessment as well as a context of use, and in section 6,

⁶¹For simplicity, I assume that the only parameters of points are moment of assessment, moment of use, moment, and history. “True” must be redefined appropriately whenever new parameters are added.

I argued that appealing to this “third grade” of truth relativity is necessary for a correct semantic account of our discourse about the future. But once we acknowledge assessment sensitivity here, the question naturally arises whether assessment sensitivity might have a use elsewhere in the semantics of natural languages. In what follows I want to explore two further applications of assessment sensitivity, to Lewis’s theory of accommodation and to evaluative relativism.

In “Scorekeeping in a Language Game,”⁶² David Lewis compares conversation to a game in which various parameters of score are adjusted in light of the players’ moves. What is distinctive about the dynamics of *linguistic* score, Lewis claims, is that the players tend to adjust the parameters in such a way that their interlocutors’ moves count as correct play. Thus, incorrect moves precipitate a change in score that makes them count as correct, a process Lewis calls *accommodation*. In the case that concerns us, the relevant parameter of score is the *standards of precision* which determine whether a sentence is true enough to be correctly asserted:

Taking standards of precision as a component of conversational score, we once more find a rule of accommodation at work. One way to change the standards is to say something that would be unacceptable if the standards remained unchanged. If you say “Italy is boot-shaped” and get away with it, low standards are required and the standards fall if need be; thereafter “France is hexagonal” is true enough. But if you deny that Italy is boot-shaped, pointing out the differences, what you have said requires high standards under which “France is hexagonal” is far from true enough.⁶³

As Lewis describes it here, accommodation is a pragmatic phenomenon: variations in the standards of precision have no effect on the actual truth conditions of sentences in context; they influence only “acceptability.” But in later papers, Lewis abandons his talk of “true enough” and partial truth and says simply that the standards of precision determine a threshold for (plain) truth:

Truth-in-English has been achieved [for the sentence ‘France is hexagonal’] if the last thing said before was ‘Italy is sort of boot-shaped.’ but not if the last thing said before was ‘Shapes in geometry are ever so much simpler than shapes in geography’.⁶⁴

⁶²Reprinted in Lewis’s *Philosophical Papers*, vol. 1, 232–49.

⁶³“Scorekeeping,” 245. See also “General Semantics,” 228–9.

⁶⁴“Index, Context, and Content,” 24.

So understood, accommodation is a semantic phenomenon.⁶⁵ We can model it by

- including in our points of evaluation a “standard of precision” parameter s ,
- invoking this parameter in the semantic clauses for particular predicates, so that their extensions at a point depend not just on the world and time parameters of that point, but also on the standard of precision parameter,
- invoking this parameter in the semantic clauses for operators like “definitely” and “sort of,” which shift the standards of precision, and
- invoking a feature of context—the contextually salient standards of precision—to initialize this parameter in the postsemantics.

The resulting semantics is capable of capturing the sensitivity of certain predicates to standards of precision that get adjusted dynamically in conversation.⁶⁶

I have intentionally been vague about the postsemantics. The standards of precision parameter is to be initialized by a feature of context: but *which* context, context of use or context of assessment? Lewis assumes that it is the context of *use*. He *has to* assume this, because his framework has no place for a context of assessment.⁶⁷ I now want to argue that the standards of precision parameter cannot be fixed by the context of use alone; the context of assessment is also needed. If this is right, then Lewis’s theory implies that “France is hexagonal” and many other sentences are assessment-sensitive. My aim is not to argue for a semantic treatment of accommodation—it

⁶⁵I am not sure why Lewis changed from a pragmatic to a semantic treatment of loose talk. Lewis’s view that speakers of a language L have a convention of truthfulness and trust in L (“Languages and Language,” reprinted in *Philosophical Papers*, vol. 1, 163–88, at 167) would seem to fit better with a semantic treatment, on which sentences like “France is hexagonal” can count as plain true. But “Languages and Language” predates “Scorekeeping in a Language Game” by four years.

⁶⁶Though I will talk of standards being “high” or “low,” we need not think of the standards of precision parameter as having simple *scalar* values. It may happen, for example, that we need to be sloppy in our application of shape terms, and at the same time very precise in our application of color terms. We can allow for this if we let the standards of precision parameter take multidimensional vectors as values. But these complexities are not important for the main semantic point I want to make here.

⁶⁷See also “Scorekeeping,” 245–6, where Lewis draws an analogy between the relativity of “The pavement is flat” to standards of precision and the relativity of “It is morning” to the time of use.

may turn out that the phenomena with which Lewis is concerned are better treated pragmatically—but to show that a semantic treatment along Lewis’s lines presupposes assessment sensitivity.

I will give two arguments for this claim. The first depends (like the arguments in the previous section) on the assumption that semantics ought to allow for objective indeterminism. The second does not depend on this assumption. Both arguments show that a single utterance can have different truth values in relation to two different contexts of assessment. This is obviously a sufficient condition for assessment sensitivity.

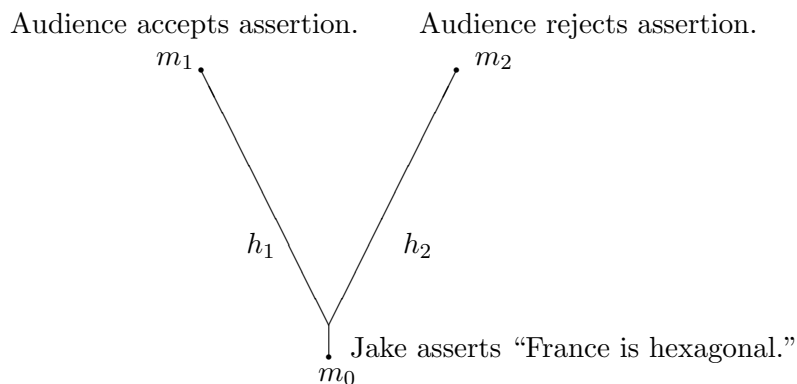
First argument

Recall that on Lewis’s view, accommodation depends to some extent on one’s *audience* or interlocutor. In order to change the standards of precision, it is not enough just to assert “Italy is boot-shaped”: you must also “get away with it”—that is, your audience must accept the assertion.⁶⁸ Only when that happens are the standards changed. But the whole point of changing the standards is to make the assertion come out *true*. If that is to be possible, the sentence asserted must be evaluated with respect to the standards of precision that are in place *after* accommodation occurs, not the standards that were in place when the assertion was made.

Now consider a scenario in which it is objectively indeterminate whether the audience will accommodate (see Figure 4). At moment m_0 , Jake asserts the sentence “France is hexagonal.” On one of the histories through m_0 (h_1), the audience accepts his assertion (at m_1), and standards of precision are lowered sufficiently for it to count as true. On another history through m_0 (h_2), the audience rejects his assertion (at m_2), and the standards remain high. Put yourself in the shoes of someone assessing Jake’s assertion from m_1 . The standards of precision have been lowered, so Jake’s assertion counts as true. Now put yourself in the shoes of someone assessing Jake’s assertion from m_2 . The standards have *not* been lowered, so Jake’s assertion counts as untrue. Thus, if Lewis is right about how accommodation works, it seems that the truth value of Jake’s assertion can depend on the context from which it is assessed. This is sufficient to show that the sentence asserted is assessment-sensitive.

⁶⁸“Scorekeeping,” 245. See also 240: “...presupposition evolves according to a rule of accommodation specifying that any presuppositions that are required by what is said straightaway come into existence, *provided that nobody objects*” (emphasis added).

Figure 4: Accommodation with indeterminism.



Second argument

The first argument depends crucially on the assumption that there is objective indeterminism, or at least that our semantics should not assume that there isn't. The second argument does not depend on this assumption: it is compatible with the view (Lewis's own) that the context of use picks out a unique future history. For the purposes of this argument, then, I will dispense with "branching histories" and use a "linear time" framework.

Suppose again that Jake asserts "France is hexagonal." Let t_0 be the time of use, and suppose that at t_1 , the audience accepts Jake's assertion and the standards of precision are lowered through accommodation. I claim that the proponent of a semantic account of accommodation ought to say both that

1. Jake's assertion is true as assessed at t_1 , and that
2. Jake's assertion is not true as assessed at t_0 .

For the whole point of accommodation is to adjust the conversational score in such a way that an assertion that would have counted as untrue (without the adjustment) counts as true.⁶⁹ Until the adjustment is made, then, Jake's assertion counts as untrue. Thereafter, it counts as true. This much relativity to a context of assessment seems built into the very idea of accommodation. If the standards of precision relative to which an assertion should

⁶⁹As Lewis puts it: "...conversational score does tend to evolve in such a way as is required in order to make whatever occurs count as correct play" ("Scorekeeping," 240).

be assessed were fully determined by the context of use, then nothing the audience did could change the truth value of the assertion.

Retroactive context change?

I find these simple arguments persuasive, but it is worth exploring how one might object to them. I can think of three general strategies (all of which I have encountered in discussing these arguments with others). The first is to suppose that when the audience accepts an assertion, triggering accommodation, the standards of precision in force at the context of use are *retroactively* adjusted. This strategy is in fact suggested by Lewis’s “general scheme for rules of accommodation”:⁷⁰

If at time t something is said that requires component s_n of conversational score to have a value in the range r if what is said is to be true, or otherwise acceptable; and if s_n does not have a value in the range r just before t ; and if such-and-such further conditions hold; then at t the score-component s_n takes some value in the range r .

Notice that it is the score-component *at* t (the time of use) that changes through accommodation. As stated, this rule makes it sound as if accommodation is an automatic process that takes effect as soon as the utterance is made. But elsewhere Lewis is clear that accommodation only happens “provided that nobody objects” to the utterance that triggers it. Since an objection (or the absence of one) would have to take place *after* the time t at which the utterance is made, Lewis must envision *retroactive* adjustment of features of the context of use.

However, this retroactive adjustment proposal, if coherently worked out, amounts to an endorsement of assessment sensitivity. Let $Standards(t)$ = the standards of precision in effect at time t . To say that the standards in effect at the time of use (t_u) *change* is to say that $Standards(t_u)$ has one value at one time, say t_1 , and another value at another time, say t_2 . But this is just to say that $Standards^*(t_u, t_1) \neq Standards^*(t_u, t_2)$, where $Standards^*(u, v)$ = the standards relevant at v for assessing an utterance that occurred at u . Although the *first* argument place in this function gets filled in by the time of use, the second one is filled in by the time of *assessment*. So any sentence whose truth value depends on the standards of precision parameter will be assessment-sensitive.

⁷⁰ “Scorekeeping,” 240.

“Fat” contexts of use?

Another strategy for resisting the arguments above is to “fatten up” contexts of use so that they *include* the audience’s response. One way to do this would be to think of an utterance as a joint action involving two agents, an utterer and an audience, just as playing tennis is a joint action. The utterance would not be over until the audience had accepted or rejected the utterer’s assertion.⁷¹ We could then say that the score parameter at a given context of use depends on the audience response, without positing *retroactive* context change.

I think that there is good reason to reject this proposal. It often takes a great deal of time for an audience to accept or reject an assertion. Think of conversations carried out in the letters columns of scientific journals, for example. It seems implausible to smear out utterances over such a great span of time. Indeed, sometimes the audience fails to respond at all (or there is no audience). Should we then say that there is no complete speech act?

These odd consequences might be worth swallowing if there were independent reasons for thinking that *all* linguistic context sensitivity must be use sensitivity. But once we have accepted *one* case of assessment sensitivity, we can no longer maintain that all context sensitivity is use sensitivity. And even if “fat contexts” allow us to avoid the conclusion that sentences like “France is hexagonal” are assessment-sensitive, they will not help in the case of future contingents. Surely it would be crazy to say that an utterance is not complete until history settles its truth value! To be sure, there might turn out to be *specific* reasons for denying that “France is hexagonal” is assessment-sensitive: reasons that go beyond an across-the-board opposition to assessment sensitivity. In that case, the “fat contexts” gambit might be worth a try. But as a strategy for avoiding assessment sensitivity altogether, it is too little, too late.

Reinterpretation

A third strategy is to say that what we do in accommodation is *reinterpret* the speaker’s original statement. For example, we reinterpret Jake’s word “hexagonal” so that it has the sense *roughly hexagonal*.

It is instructive to compare what happens here with ordinary cases of

⁷¹For the idea of uttering as a joint action, see Belnap, “Double Time References,” 7. On Belnap’s view, however, the audience’s role is simply to “lend an ear;” the utterance does not include the audience’s *response*.

reinterpretation. Suppose Sam says “I’m going to the bank,” and we assume he’s headed off for some fishing, but the next thing he says is “I’ve got to check on that loan I applied for.” We may then reinterpret his utterance, taking “bank” now in the sense of “financial institution” rather than “edge of river.” But in doing this, we concede that we *made a mistake* in our original construal of his utterance. We would have been *wrong* to assess his assertion for truth or falsity in light of our original interpretation. And if we *hadn’t* reinterpreted, we would have *misunderstood* Sam’s claim. There is no analogy to this in accommodation. If the audience chooses to reject Jake’s assertion, then it is perfectly correct to construe it in the strict sense and take it to be false.⁷²

One might still persist in talking of reinterpretation, adding that this talk carries no implication that one’s original interpretation was in error. But now it is not clear that the talk of reinterpretation is a way of *avoiding* assessment sensitivity, rather than *re-describing* it. Reinterpretation here can’t be a matter of assigning a different lexical meaning to one of Jake’s words, or assigning a different syntactic structure to one of his sentences, or taking an ordinary indexical like “I” or “this” to have a different referent: these would all require admitting a mistake. It can only mean taking “hexagonal” to be governed by a different standard of precision, a standard determined by features of the context of assessment as well as the context of use. And doesn’t this amount to accepting that “hexagonal” is assessment-sensitive after all?

Moreover, talk of reinterpretation would be quite unnatural in the future contingents case. Do we really want to say that when the sea battle takes place, we *reinterpret* yesterday’s utterances of “There will be a sea battle tomorrow?” This sounds like something only someone in the grip of a theory would say. Even if we did say this, it’s not clear it would be incompatible with acknowledging the assessment sensitivity of future contingents, for the reasons given above. What is the difference (other than verbally) between acknowledging a kind of reinterpretation that is specifically allowed by language and doesn’t carry any implication that the original interpretation was in error, and acknowledging assessment sensitivity?

⁷²If you disagree, then you are rejecting Lewis’s semantic treatment of accommodation. Again, it is not my project here to defend that theory, but only to show that it requires acknowledgement of assessment sensitivity.

Postsemantics for accommodation

Let us grant, then, that accommodation *does* involve assessment sensitivity, and return to the question we raised above: how do we treat the standards of precision parameter in the postsemantics? I have argued that the treatment of this parameter will depend somehow on the context of assessment, but I have not said how.

The simplest possibility is to set the standards of precision parameter to the standards prevailing at the context of assessment. But this won't work. To see why, return to our example (Figure 4). We get reasonable results if the context of assessment is at m_0 , m_1 , or m_2 . But consider a moment $m_3 > m_1$ such that between m_1 and m_3 the standards of precision are raised again. Do we really want to say that Jake's assertion at m_0 is untrue, as assessed at m_3 ? Surely it would be more reasonable—and more faithful to Lewis's proposal—to say that only the standards prevailing *immediately after* the audience's response to Jake's assertion are relevant for assessing this assertion.

How can we capture this more reasonable proposal in the postsemantics? Somehow we need to relativize the standards of precision to *utterances*, so that at a single context of assessment, different standards of precision can be used to evaluate different utterances. Since a context of use is a concrete utterance episode, we can relativize the standards of precision to contexts of use rather than utterances themselves. Let *Standards* be a function that takes as arguments a context of use and a context of assessment and delivers as output a standard of precision:

Definition of *Standards* function: For all contexts of use U and contexts of assessment A , $Standards(U, A) =$

- the standards of precision in place at U , if at A the audience has neither accepted nor rejected the assertion made at U
- otherwise, the standards of precision in place immediately following the audience's acceptance or rejection of the assertion made at U .

Then the following postsemantics will work:

Postsemantics for accommodation: A sentence Φ is true [false] at a context of use U and context of assessment A iff Φ is true [false] at every point $\langle D, s, a \rangle$ such that

- $D =$ the domain intended by the speaker at U

- $s = Standards(U, A)$
- $a =$ an assignment of objects from the domain to the variables.

If Lewis is right that accommodation is a semantic phenomenon, and if this is the correct account of it, then a mild sort of assessment-sensitivity turns out to be very common. For many expressions of natural language can be applied with different standards of precision.

8 Evaluative relativism

Since antiquity, philosophers have proposed that the truth of claims about matters of morality, aesthetics, taste, or fashion are “relative.” I think that the framework I have been articulating in this paper can help to clarify the *content* of such proposals, illuminate what is attractive about them, and show what would be required to argue for them. For the sake of concreteness, I will focus on relativism about beauty.

One way to think of relativism about beauty is as a claim about the semantics of “beautiful.” The truth of sentences containing “beautiful” is relativized to some parameter—say, an “aesthetic standard”—so that a sentence like “Tiny feet are beautiful” can be true on one aesthetic standard, false on another. This is as far as many treatments of relativism get. But we are now equipped to ask: what *kind* of relativization of truth is this?

On any version of the proposal, we will need at least the first grade of truth relativity: an “aesthetic standard” parameter in our points of evaluation. Perhaps this parameter can even be *shifted*, by operators like “to the nineteenth-century Chinese . . .” But by itself, the relativization of truth to this parameter is not enough for an interesting version of value relativism. After all, the same kind of relativization is required for the semantics of *any* moderately complex language. No one called Tarski a relativist for saying that sentences are true or false relative to a satisfaction sequence, or Kripke for saying that sentences are true or false relative to a possible world!

The interesting question is how the aesthetic standard parameter gets treated in the postsemantics. Is it determined by the context of use (like the domain parameter)? Constrained by the context of use (like the assignment parameter)? Determined by the context of assessment (like the moment of assessment parameter)? Constrained by the context of assessment (like the history parameter)? Or determined jointly by the context of use and the context of assessment (like the standards of precision parameter)?

If the aesthetic standard parameter is determined (or constrained) by the context of use, then sentences containing “beautiful” will be use-sensitive, just like sentences containing ordinary indexicals or tensed verbs. On this construal, there can be no worries about the *coherence* of relativism. The truth of “Tiny feet are beautiful” is relative in much the same way that the truth of “Everyone is here” is relative. It exhibits the second grade of truth relativity.⁷³

This indexical version of relativism allows us to count *both* Hao’s assertion of “Tiny feet are beautiful” *and* Marco’s assertion of “Tiny feet are not beautiful” as true (assuming that the respective contexts of use pick out different aesthetic standards).⁷⁴ It thereby achieves one of the main goals of relativism: avoiding having to say that one of the parties to this disagreement is *making a mistake*.⁷⁵ But it does so by denying, in effect, that there is any disagreement here at all. On the indexical account, Hao’s assertion no more contradicts Marco’s than my assertion of “I was born in New Mexico” contradicts your assertion of “I was not born in New Mexico.” The disagreement is merely apparent, merely verbal.⁷⁶

Is there any way to acknowledge the disagreement between Hao and Marco as genuine, yet still regard it as “blameless”? Not in traditional semantic frameworks (which give no semantic role to the context of assessment). But if we admit the third grade of truth relativity—relativity to a context of assessment—another form of relativism becomes possible. On this form of relativism, the aesthetic standard parameter is determined (or constrained) by the context of *assessment*, rather than the context of use. Suppose, for concreteness, that it is initialized by the aesthetic standards of the person who is assessing the utterance.⁷⁷ Then we can say the following:

⁷³Gilbert Harman’s version of moral relativism fits this mold: “Moral Relativism Defended,” *Philosophical Review* 84 (1975), 3–22.

⁷⁴I assume that these are genuine assertions, and not, say, expressions of attitudes. Expressivism faces substantial difficulties: for a recent survey, see chapter 4 of Max Kölbel, *Truth Without Objectivity* (London: Routledge, 2002).

⁷⁵I will not here concern myself with whether this motivation is sound, or whether relativism is the only way to meet it. For useful discussion, see Kölbel, chapter 2.

⁷⁶For this kind of criticism, see Philip Hugly and Charles Sayward, “Moral Relativism and Deontic Logic,” *Synthese* 85 (1990), 139–152, at 142–3.

⁷⁷Other options are possible. For example, one might quantify over the aesthetic standards of every person present when the utterance is being assessed. In this case, the aesthetic parameter would be constrained, but not determined, by the context of assessment. Or, one might quantify over the aesthetic standards of the assessor *and* the utterer; this would be a mixed case, like that of the standard of precision parameter (section 7, above).

- Hao’s assertion is true as assessed by Hao.
- Hao’s assertion is false as assessed by Marco.
- Marco’s assertion is true as assessed by Marco.
- Marco’s assertion is false as assessed by Hao.

We can now do justice *both* to the intuition that Hao and Marco are genuinely (and not just verbally) disagreeing, *and* to the intuition that their disagreement is blameless. They are genuinely disagreeing, because relative to any one context of assessment, only one of their assertions can be true. (Indeed, the sentences they assert are logically inconsistent, in the sense that they cannot both be true relative to *any* context of utterance and context of assessment.) Nonetheless, the disagreement between Hao and Marco is blameless, because in the face of a challenge, each of them can produce unimpeachable grounds for the truth of his assertion, relative to *his own* context of assessment. So neither can be obliged by the others’ challenges to withdraw his assertion.

My point here is not to endorse any particular kind of relativism about beauty. All I have sought to do is lay the semantic groundwork for the formulation and discussion of relativist views. Too many discussions of relativism say merely that the truth of sentences is relative to some parameter, without being explicit about how that parameter is to be treated in the postsemantics.⁷⁸ My goal has been to show how many structurally distinct possibilities there are and to provide a framework within which a determinate version of relativism can be stated. Once a version of relativism has been fully stated—with both a compositional semantics and a postsemantics—it will be clear exactly what substantive work is necessary to defend it. If one accepts the account of assertion sketched in section 5, above, then justifying a version of relativism will require showing that the norms governing withdrawal of assertions conform to the predictions yielded by the postsemantics and (Assertion-3).

Evaluative relativism has always had to struggle with doubts about its coherence. Harman’s strategy for quieting these doubts was to assimilate his moral relativism to ordinary indexicality, the coherence of which no one

⁷⁸It does not help to speak of propositions rather than sentences. For example, Kölbel characterizes “non-tame” relativism as any view of the form: “For any x that is a proposition of a certain kind K , it is relative to [some parameter] P whether x is true” (*Truth Without Objectivity*, 119). But this cannot suffice: virtually all standard views of propositions hold that contingent propositions have different truth values relative to different possible worlds, and they are not on that account relativist!

could deny.⁷⁹ But there remained doubts about the coherence of forms of value relativism that did not follow this pattern—those that let truth depend on facts about the *assessor* of a claim, or the context of assessment. I hope to have gone some way toward removing these doubts, by showing not only that assessment sensitivity is coherent, but that much of our *non-evaluative* discourse is assessment-sensitive. If we need assessment sensitivity to make sense of accommodation, or talk about the future, then we can hardly make *formal* objections to assessment-sensitive accounts of moral or aesthetic language. If these forms of relativism are false, that must be shown through substantive consideration of our practices of using moral and aesthetic language.

9 Conclusion

I have argued that the truth value of an assertion can depend not only on the sentence used and the context in which it is asserted, but on the context in which it is *assessed*. I have argued that this much relativism about truth is not only coherent, but *required* if we are to have a satisfactory semantics for future contingents. And I have sketched a couple of further applications of this “third grade of truth relativity.”

Philosophers wedded to orthodox frameworks will no doubt be able to produce epicycles that allow them to handle my examples without acknowledging assessment sensitivity. I welcome attempts to do this, and I recognize that there is always a presumption in favor of old ways of doing things. But I think that the framework advocated here will ultimately provide the most elegant and satisfying account of the data. Notice also that it is a *generalization* of the orthodox frameworks, which emerge as a special case when the language contains no assessment sensitivity. We can therefore explain the theoretical success of the old framework by pointing to the comparative rarity of assessment-sensitive expressions. (Future contingents were always a stumbling block in the orthodox framework; now we can see why.) As more and more applications for assessment sensitivity are found—and I believe there are important applications to the semantics of epistemic terms like “might” and “know” and to the semantics of past scientific discourse—I believe that the framework I have advocated here will look more and more attractive.

⁷⁹ “Moral Relativism Defended.”