Truth and Subjectivity[*]

John MacFarlane[†]

April 2, 2007

My aim in this talk is to describe and motivate the framework I've been developing for bringing relative truth—or, as I prefer to call it, "assessment sensitivity"—into the study of language and communication. Those of you who are familiar with some of my other writings on the subject will have heard much of this before, but there is at least one new twist to the way I'll do things here: I'm going to talk only of propositions—the things we assert, believe, and ponder—and not say anything about the meanings of the linguistic expressions we typically use to express these propositions. I am impressed by the fact that which sentence one uses to assert a given proposition on a particular occasion will depend on a huge number of contextual factors, including one's aims in expressing it, one's views about the actual state of the world, and one's expectations about the intelligence, background assumptions, and linguistic competence of one's interlocutors. Although for various simplificatory purposes it is useful to indulge in the fantasy that a systematic theory of meaning can associate propositional contents with sentence/context pairs in a straightforward, rule-governed way, I would not want to stake any money on this, and I would especially not want to stake my research program on it. Thus, I will talk today only of sayables and thinkables, and not of their sentential vehicles.

1 Subjective discourse

One historically significant motivation for appealing to relative truth is the desire to understand the sense in which certain claims we make, and certain thoughts we entertain, are *subjective*.

I'll call a claim *subjective* if its truth depends not just on how things are with the things it is about, but on how things are with some subject who is not part of its subject-matter. An *objective* claim is one that is not subjective. In talking of "claims" here, I mean to be talking of the *contents* of claims— propositions—not acts of claiming. I don't intend to get much mileage out of these definitions—after all, there is at least as much unclarity about the notion of "aboutness" as there is about subjectivity—but it should suffice for our purposes here.[1]

In the sense that concerns me,

(1)   Joe believes that Porsches are sexier than Ferraris

is an objective claim, not a subjective claim. Its truth depends only on how things are with the things it is about (Joe and his mental state). However, the thing Joe is said to believe,

(2)   that Porsches are sexier than Ferraris

is, plausibly, subjective. For its truth may seem to depend not just on how things are with Porsches and Ferraris, but on the reactions of the subject who is considering the claim, or perhaps the subject who has made the claim. Different people find different characteristics sexy, after all, and it is not at all obvious that we should explain this by pointing to difficulties in *determining* whether something is really sexy. The disagreement here isn't much like disagreements about, say, the age of the earth.

Here are some more examples of claims that have been thought, by at least some philosophers, to be subjective:

(3)   Oyster mushrooms are *delicious*.

---

[1]Note that if the claim quantifies over a domain, then we will have to count the whole domain as part of its subject-matter. Otherwise, the claim that Joe is standing next to someone who likes flowers would count as subjective.

(4) That joke was *funny.*

(5) His humor is *obscene.*

(6) It's *cruel* to amputate a thief's hand.

(7) Goldbach's Conjecture *might* be false.

(8) It's *likely* that Joe and Sue are on the 6:10 train.

(9) *If* the valve is open, water is coming out of spillway #1.

Some of these are more evidently subjective than others. You may accept that claims about what is "delicious" are subjective, for example, while denying the subjectivity of discourse about what is "cruel." For my purposes here, I only need you to agree that some of these examples are examples of "subjective" discourse, in the sense I've described.

Could one take the hard line that they're *all* objective? There are two ways in which one might do this. The first—*hard-core objectivism*—is to deny that the truth of the relevant claims depends at all on how things are with subjects not explicitly mentioned in the sentences listed above. I can imagine someone being a hard-core objectivist about, say, cruelty, or knowledge. But to take this position across the board seems impossible.

Consider (7), for example. This is an expression of epistemic possibility, of course. The claim isn't that Goldbach's Conjecture *could have been* false: since mathematical truths are presumably necessary, that would amount to saying that it *is* false, and we don't know that. Rather, the force of the claim is something like this: given what is known, we can't rule out the falsity of Goldbach's Conjecture. But "what is known" by whom? Here evaluating the truth of the claim requires specifying a subject (or subjects)—perhaps the speaker or thinker, or the members of some group suitably related to the speaker or thinker. You could try to make "what is known" non-relative, I suppose, by understanding it as: what will be known by the community of knowers "at the end of inquiry." But if that's the force of the claim, we're in no position to make it—yet we do make it.

Similarly with "Oyster mushrooms are delicious." We generally take ourselves to be justified in thinking or asserting that something is delicious solely on the basis of our own affective reactions. Yet we're well aware that others' affective reactions differ from our own. If the truth of claims of deliciousness does not depend at all on our own affective reactions, we're being

remarkably chauvinistic and unreflective. Here hard-core objectivism looks very unattractive.

The second way to defend the objectivity of our apparently subjective discourse is to construe these claims as claims *about* the subjects on whom their truth depends. As we've already noted, the claim that

(10) Joe likes the taste of Oyster mushrooms.

is perfectly objective. So, if the proposition expressed by (3) (as asserted by Joe) is just (10), its subjectivity is merely apparent. Similarly, if (7) is a compressed way of expressing what could also be expressed by

(11) I do not know that Goldbach's Conjecture is false,

or even

(12) Mathematicians in March 2007 do not know that Goldbach's
      Conjecture is false,

then it is, despite appearances, fully objective. I call this strategy for explaining away apparent subjectivity the "subject as subject-matter" strategy. It is, I think, the most common approach, much more common than hard-core objectivism, so I will also call it the "standard" strategy.

In considering this strategy, I'll abstract from questions about just how the unexpressed subject gets into the content of the claim. One might, for example, take "delicious" to be indexical, like "here," or to be associated with an unpronounced demonstrative element in the syntax. Or you might appeal to some looser, less rule-governed process of "free enrichment." It really doesn't matter for our purposes, just as long as we end up with a proposition that is about the relevant subject.

2 Lost disagreement

There is, I think, a very basic problem with this strategy: it can't explain the *disagreement* we perceive in subjective domains.

Here are some examples of conversations that should seem relatively natural:

1. A: Oyster mushrooms are delicious.
   B: No, they aren't. They're yucky.

2. A: Joe might be in Boston.
   B: No, he can't be in Boston; I just saw him down the hall.

Both cases appear to be cases of disagreement. B thinks A has *gotten things wrong*. But disagreement would not be expected here, on the "subject as subject-matter" analysis:

1. A has asserted that oyster mushrooms taste good to him. B is now denying this (!), apparently on the ground that they don't taste good to *him* (B). But this is no ground at all. If this analysis is correct, we should expect B to be able to say, "You're right, I'm sure. But they're yucky."

2. A has asserted that what A knows is consistent with Joe's being in Boston. B should just acknowledge this as true—it's not contradicted by the fact that what B knows *does* rule out Joe's being in Boston. Consider how weird it would be for B to say, "You're right, I'm sure. But I just saw him down the hall."

There's an obvious strategy for fixing the standard analysis: move from an individual to a *group* that includes both speakers. For example, in (1) we might say that A is asserting (and B denying) that oyster mushrooms taste good to them both (or, to those like them in some relevant respect). Similarly, in (2), we might say A is asserting that what is known by a contextually relevant group, including both A and B, does not rule out Joe's being in Boston.

But as I've argued at length elsewhere,[2] this won't work. First, once you start dealing with perceptions of disagreement by widening the community of subjects you're incorporating into the claim's subject-matter, there's no good place to get off the bus. After all, hidden eavesdroppers (even spatially and temporally remote ones, like future historians) will behave the same as B in the dialogues above. So we'd really have to go to a *universal* community, which is absurd.

---

[2]For the general argument, see "Relativism and Disagreement," *Philosophical Studies* 132 (2007), 17-31. For the specific cases of epistemic modals and knowledge attributions, see "The Assessment Sensitivity of Knowledge Attributions," *Oxford Studies in Epistemology* 1 (2005), 197-233, and "Epistemic Modals are Assessment-Sensitive," forthcoming.

Second, the more we widen the community, the harder it becomes to understand how A could have taken herself to be warranted in making the claim in the first place. People shouldn't, on the basis of their *own* affective reactions or their own information, make claims about the reactions or information of large groups.

Now, I need to qualify this a bit. I'm not saying that a speaker can *never* use "Oyster mushrooms are delicious" to assert that oyster mushrooms taste good to him, or to members of some particular group. Nor am I saying that a speaker can *never* use "Joe might be in Boston" to assert that what *she* knows does not rule out Joe's being in Boston. As Donnellan reminds us in his discussion of Humpty Dumpty, just about any sentence can, in the right setting, be used to assert just about anything.[3] What I will insist on is that we *sometimes* assert and believe subjective contents whose subjectivity can't be well understood using the "subject as subject-matter" strategy. (In fact, I think this is the default, but I needn't argue that here.)

3   Relativizing proposition truth

If the arguments of the previous section are correct, then we won't be able to use the "subject as subject-matter" strategy to explain away all cases of apparently subjective discourse. Instead of arguing that apparently subjective claims are really objective, we'll have to explain what their subjectivity consists in. How can the truth of a claim depend on how things are with something other than its subject matter?

Here it might be useful to consider an analogy. The standard view of present-tensed sentences is that they express propositions that are about a particular time—normally, the time at which they are uttered. So, if I now say "Socrates is sitting," I am expressing the proposition that Socrates is sitting at noon on March 8, 2007. On this view, the time of utterance gets to be part of the *content* of my claim—part of what my claim is *about*.

But some philosophers—*temporalists*—reject this view. According to temporalists, whenever I say, "Socrates is sitting," I express the same *time-neutral* or "tensed" proposition, the proposition that Socrates is sitting. This proposition has different truth values relative to different times of evaluation, just as standard "eternalist" propositions are thought to have differ-

---

[3] "Putting Humpty Dumpty Together Again," *The Philosophical Review* 77 (1968), 203-215, at 213.

ent truth values relative to different worlds of evaluation. And, just as the *world* of utterance is not usually thought to be part of the content expressed by contingent sentences, on the temporalist view, the *time* of utterance is not part of the content. My assertion may *concern* a particular time, the time at which it is made, but what I've asserted is not *about* any particular time.[4]

So here is a model for us to follow. On the temporalist view, the truth of a tensed claim depends on a time, not because the claim has a time as its subject-matter, but because the claim's content has truth values only relative to times. By relativizing the truth of propositions, then, we might explain how discorse about what is delicious, or what might be the case, can depend for its truth on a standard of taste or a body of information, yet without being *about* these things.

[Aside: by "proposition," I just mean the contents of the propositional attitudes. I want to be completely neutral here about whether these are structured or unstructured, fine-grained or coarse-grained, etc. These questions are independent of the issue I want to focus on.]

Now, this step in itself—relativizing propositional truth—is by no means a radical step. Kaplan, in "Demonstratives," countenances contents that have truth values relative to worlds and times.[5] Indeed, he's prepared in principle to go further, if a proper semantic analysis demands it. He calls that to which the truth of contents must be relativized a "circumstance of evaluation":

> By ["circumstances"] I mean both actual and counterfactual situations with respect to which it is appropriate to ask for the extensions of a given well-formed expression. A circumstance will usually include a possible state or history of the world, a time, and perhaps other features as well. The amount of information we require from a circumstance is linked to the degree of specificity of contents, and thus to the kinds of operators in the language. (502)

---

[4]I take the terminological distinction between "about" and "concerns" from John Perry. See, for example, his "Thought Without Representation," *Proceedings of the Aristotelian Society*, s.v. 60 (1986), 137–52.

[5]"Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals," in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press, 1989), 481–566.

> What sorts of intensional operators to admit seems to me largely
> a matter of language engineering. It is a question of which fea-
> tures of what we intuitively think of as possible circumstances
> can be sufficiently well defined and isolated. If we wish to isolate
> location and regard it as a feature of possible circumstances we
> can introduce locational operators: 'Two miles north it is the
> case that', etc. ... However, to make such operators interesting
> we must have contents which are locationally netural. That is,
> it must be appropriate to ask if *what is said* would be true in
> Pakistan. (For example, 'It is raining' seems to be locationally
> as well as temporally and modally neutral.) (504)

If we can contenance contents that are world-neutral, time-neutral, or even
location-neutral, then there seems no conceptual obstacle to investigation
contents that are *taste-neutral*. Assuming we still have world neutrality, our
basic relativized propositional truth predicate would be

p is true at $\langle w, s \rangle$,

where s represents a standard of taste. While we're at it, why not contents
that are taste- and epistemic-state- neutral?

p is true at $\langle w, s, e \rangle$

Or, individual- and time-neutral?

p is true at $\langle w, t, i \rangle$

(You'll recognize these latter as Lewisian "centered-worlds" propositions,
which Andy Egan has been trying to press into service to handle "subjective
discourse." I think Egan's approach is problematic, but won't get into the
reasons here.[6])

I don't see any good *general* argument against this kind of relativization
of propositional truth. The fate of the temporalist's relativization to *time*,
which has fallen out of favor of late, is instructive. People give two reasons
for eschewing the relativization to times:

---

[6]The curious can see "Epistemic Modals are Assessment-Sensitive," fn. 26.

1. Tenses in natural languages are better understood in terms of quantification and anaphora than as operators. This undercuts Kaplan's reason for relativizing to times.

2. Considerations of same-saying and same-believing seem to suggest that the time of utterance (unlike the world of utterance) is part of propositional content. For example, if I say now "It's sunny," and you say tommorow "It's sunny," and we're both sincere, it doesn't seem to follow that there's something that you and I both asserted, and that you and I both believe.

I won't pause here to assess the merits of these arguments.[7] What I want to point out is that neither kind of consideration poses a problem for "taste-neutral" propositions. (1) does not, because our motivation for countenancing "taste-neutral" contents was not our desire to treat anything as an operator. (2) does not, because here the samesaying and samebelieving intuitions are largely in our favor. We *do* have a strong inclination to say that when I say "oyster mushrooms are delicious" and you say "oyster mushrooms are delicious," we have asserted the same thing, and expressed the same belief, even if we have different tastes. (This is just the flip side of the point made earlier about disagreement.) So, in some ways, the case for taste-neutral contents is stronger than the case for time-neutral contents ever was.

4   Nonindexical contextualism

Lots of people stop here. But this move turns out to be *insufficient* for making sense of disagreement in subjective domains. We now have it that A accepts a "taste-neutral" proposition, **that oyster mushrooms are delicious**, which B rejects. But, contrary to what you might have thought, this isn't sufficient for A and B to genuinely disagree.

The general point here can be illustrated without considering funny taste-neutral propositions:

1. A (in w_1) asserts that Dodos are not extinct in 2006.
   B (in w_2) denies this same proposition.

---

[7] I'm skeptical about the second one in particular, for reasons articulated by David Lewis, "Index, Context, and Content," in *Papers in Philosophical Logic* (Cambridge: Cambridge University Press, 1998), at 41.

2. A (at noon) asserts (the time-neutral proposition) that it is daytime.
B (at midnight) denies this same (time-neutral) proposition.

When there is genuine disagreement, I take it, it cannot be that each party regards the other's acceptance or rejection as objectively correct, accurate, or "true in its context." But in (1), both A and B regard the other's acceptance or rejection as objectively correct. Though the proposition at issue isn't *about* any particular world—it is "world-neutral"—each speech act *concerns* a particular world, and it is to this world that both parties look in assessing the accuracy or correctness of the assertion or denial. (The useful terminological distinction between "about" and "concerns" is due to John Perry.) Similarly, in (2), B can very well say that A was correct—not just in the sense of being justified, but in the sense of being accurate—to assert at noon what B denies at midnight.

Because the acceptance and rejection in both cases *concerns* different worlds (or times), they are *both* objectively correct, even though their contents are contradictory. And if they are both objectively correct, there is no real disagreement here.

Taking these views as our model, we could say that although **oyster mushrooms are delicious** is taste-neutral and not *about* any particular subject's standard of taste, an assertion (or acceptance) of it *concerns* the speaker's (thinker's) taste, and no one else's. On this view, there would be no real disagreement between someone who asserts **oyster mushrooms are delicious** and someone who denies this—any more than between speakers who assert and deny a single tensed content at different times. So we would still, on this view, have the problem of lost disagreement. Even though the speaker's tastes are no longer made part of the subject matter, still the speech acts (or mental acts of acceptance or rejection) *concern* the speaker's tastes. I've elsewhere called this view "nonindexical contextualism." (Francois Recanati has recently been defending what is essentially the same view under the name "moderate relativism.")

Now, some people might be attracted by a view that gives the mere appearance of disagreement in the cases at issue. Perhaps all there really *is* is the appearance of disagreement. But if we want more than that, we're going to have to say more than we have so far, since what we've said so far is all compatible with a view that doesn't secure real disagreement.

It will have to be a view on which assertions about what is delicious do not *concern* any particular subject's taste, and on which assertions about what

is epistemically possible do not concern any particular subject's (or any group's) information. We could do this by saying that it is the *assessor's* taste, or information, not that of the speaker or any group picked out by its relation to the speaker, that matters for the correctness of such assertions. But we need some new concepts to talk clearly about this.

5 Assessment sensitivity

To make these ideas more precise, I'm going to follow Kaplan in distinguishing between two different dimensions along which we might relativize truth.

1. We can talk of the truth of a content at a circumstance of evaluation (e.g. a world/time pair, or a world/standard-of-taste pair). Here we're asking: would this content be true if the world were like so, and the relevant standard of taste were such-and-such?

2. We can also talk of the truth of the content relative to a *context*. Here we're asking: could this content be truly asserted or believed in this context? (Aside: Kaplan himself only talked of the truth of *sentences* relative to contexts, but it is easy enough to extend what he says to contents.)

When circumstances of evaluation are world/time pairs, the relation between these two relativized notions of content-truth is straightforward:

p is true at c iff p is true at $\langle w_c, t_c \rangle$, where $w_c =$ the world of c and $t_c =$ the time of c.

If we used the same model for "delicious," we'd say

Nonindexical contextualism: p is true at c iff p is true at $\langle w_c, s_c \rangle$, where $w_c =$ the world of c and $s_c =$ the speaker's standard of taste at c (or, the standard of taste picked out by the speaker's intentions and the objective situation at c).

But this gives us nonindexical contextualism, which I've argued does not secure genuine disagreement. To get genuine disagreement, we need to give a

11

truth-determinative role to the context from which the content is *assessed*—
what I've called the *context of assessment*. So now we have two contexts,
and "truth as used at c_u and assessed from c_a" replaces "truth at c" as the
target notion. This is related to truth at a circumstance of evaluation in the
following way:

> **Relativism:** p is true as used at c_u and assessed from c_a iff p
> is true at $\langle w_{c_u}, s_{c_a} \rangle$, where $w_{c_u} =$ the world of c_u and $s_{c_a} =$ the
> assessor's standard of taste at c_a (or, the standard of taste picked
> out by the assessor's intentions and the objective situation at c_a).

Note that on the nonindexical contextualist account, the context of assess-
ment plays no role. On this account, the proposition that oyster mushrooms
are delicious and the proposition that oyster mushrooms aren't delicious can
both be true, relative to the contexts in which they are asserted, provided
their asserters have different tastes.

On the relativist account, this can't happen: at most one of these proposi-
tions will come out true, relative to the context in which it is used and our
context of assessment. In this way, we make room for genuine disagreement
about matters of taste, while preserving a role for subjectivity.

Of course, which one of these propositions *is* true will depend on our vantage
point—which is why this amounts to a kind of relativism about truth. In
fact, this is how I *define* relativism about truth:

> A proposition p is *assessment-sensitive* iff for some contexts
> c_1, c_2, c_3, p is true as used at c_1 and assessed from c_2 but not
> true as used at c_1 and assessed from c_3.

> To be a *relativist about truth* is to hold that there are assessment-
> sensitive propositions.

6   What about the equivalence schema?

A common objection to relativism about truth is that it somehow requires
rejecting the equivalence schema, which supports semantic ascent and de-
scent. We're now in a position to evaluate this objection.

The version of the equivalence schema that is relevant here is the

> **Propositional equivalence schema:** The proposition **that** p
> is true iff p.

The thing to notice is that the idea of disquotation makes sense only for a *monadic* truth predicate. Once *any* relativization is added, it no longer makes sense to disquote:

> The proposition that p is true at w iff ???[8]

Thus the suggestion that we relativize truth of contents to more than just worlds does not itself require us to give up the equivalence schema. (Similarly, the relativization of truth to contexts of use and contexts of assessment.)

To state the equivalence schema, the relativist needs to introduce a monadic truth predicate. This will be a predicate of the object language, not the semantic metalanguage—since the theorist is taking a point of view outside any particular context. (This is a methodological point: it's not that the theorist doesn't occupy a particular context of assessment, but that she avoids using assessment-sensitive terms in describing truth conditions, in just the same way that Kaplan avoids using indexicals in stating truth conditions for indexical sentences.) We can define a monadic truth predicate by giving its extension at every circumstance of evaluation (here I abstract from worries about paradoxes):

> **Monadic 'true':** The extension of 'true' at $\langle w, s \rangle$ is the set of propositions that are true at $\langle w, s \rangle$.

This is a completely natural semantics for monadic 'true': it's hard to see what else would be a better candidate. And it turns out that, given this clause for the extension of 'true', every instance of the propositional equivalence schema will be true at every circumstance of evaluation, and hence also at every context of use and context of assessment. (Whatever circumstance $\langle w, s \rangle$ we choose, the right and left hand sides of the biconditional will have the same truth value at $\langle w, s \rangle$. I assume here that "the proposition that p" rigidly denotes a proposition.)

---

[8]Perhaps an analogue of disquotation could be preserved here, by filling in the "???" with "in w, p," and understanding "in w" along the lines of "in Australia." See David Lewis, *On the Plurality of Worlds* (Oxford: Blackwell, 1986), 5-7. But then it is not clear why someone who relativized worlds to standards of taste could not simply use "by standard s" to the same purpose.

So this criticism falls flat. But there *is* a legitimate concern in the vicinity. Granted that our doubly relativized truth predicate is not the ordinary (monadic) truth predicate we use in ordinary speech, but a piece of *technical* vocabulary, we need to say something about how it is connected up with other parts of our theories of language and communication, so we can see the practical significance of going for a relativist semantic theory as opposed to a nonrelativist one. I want to emphasize, though, that this is a burden faced by nonrelativists, too—by anyone who uses "true at a context" in a truth-conditional semantic theory.[9]

7   Relative truth and assertion

What makes the relativist's doubly context-relative predicate a *truth* predicate, and not just some bizarre term we can define but have no reason to care about? Here I'll be very brief (see my "Making Sense of Relative Truth"[10] for a fuller discussion).

Here's a way of making this question more focused. The concept of truth plays a certain role in our cognitive and communicative lives:

- Truth is what we commit ourselves to in making assertions.

- Truth is what we aim at in belief.

Can an assessment-relative notion of truth play these roles?

Well, what is it to commit oneself to the truth of a proposition? What would count as honoring, or as violating, such a commitment? What would one have to *do* to honor or violate it? Here is an answer that seems to me plausible:

> **Commitment to truth (nonrelativist)** To commit oneself to the truth of a proposition p (at some context c) is to undertake commitments to
>
> - vindicating the assertion (that is, providing grounds for thinking that p is true-at-c) in response to a legitimate challenge.

---

[9]The point goes back at least to Dummett's classic article "Truth," *Proceedings of the Aristotelian Society* n.s. 59, 141–62.

[10]*Proceedings of the Aristotelian Society* 105 (2005), 321–39.

- retracting the assertion (withdrawing all the associated commitments) when p is shown to be untrue-at-c.

Here the (nonrelativist) semanticist's notion of "truth at a context of use" is used to help explicate "commitment to the truth of p." So someone who had a definition of "truth at a context" for arbitrary propositions could use this account to see what kinds of normative changes would be effected by the assertion of an arbitrary proposition at an arbitrary context.

I think it's the possibility of drawing a connection like this between the semanticist's notion of "true at a context of use" and the use we make of contents—e.g. in asserting them—that makes it appropriate to call the semanticist's notion a relativized notion of *truth*.

Now, what if truth is assessment-sensitive? Can we still understand commitment to truth? I think that we can. Note that there are two contexts that might be relevant to such a commitment: that in which the assertion was made and that which one currently occupies. So there's an extra degree of freedom in the notion of commitment to truth that we can appeal to:

> **Commitment to truth (relativist)** To commit oneself to the truth of a proposition p (at some context c) is to undertake commitments to
>
> - vindicating the assertion in response to a legitimate challenge (that is, providing grounds for thinking that p is true as used at c **and assessed from the context in which one is meeting the challenge**).
> - retracting the assertion (withdrawing all the associated commitments) when p is shown to be untrue as used at c **and assessed from the context one occupies in considering the retraction**.

So I think we can still make sense of commitment to truth if truth goes assessment-sensitive.

Note that this proposal allows us to generate definite normative predictions from relativist proposals, and to compare them with nonrelativist proposals. Finally, it is *conservative* in the following sense: if there are no assessment-sensitive propositions, it reduces to the nonrelativist version. For this reason, I would hope that even someone who doubts that we ever do think or

claim anything assessment-sensitive could accept this account as a neutral framework for adjudicating and investigating claims that particular areas of discourse involve assessment sensitivity.