

XIV*—MAKING SENSE OF RELATIVE TRUTH

by John MacFarlane

ABSTRACT The goal of this paper is to make sense of relativism about truth. There are two key ideas. (1) To be a relativist about truth is to allow that a sentence or proposition might be assessment-sensitive: that is, its truth value might vary with the *context of assessment* as well as the context of use. (2) Making sense of relativism is a matter of understanding what it would be to commit oneself to the truth of an assessment-sensitive sentence or proposition.

Analytic philosophers tend to regard relativism about truth as hopelessly confused, easily refuted, and even a sign of deficient intellectual character. This attitude is not entirely unreasonable. Proponents of relativism have focused much more on *motivating* their doctrine than on making sense of it, or even stating clearly what it is to be a relativist about truth. But if relativists have underestimated the difficulties here, their opponents have overestimated them. It is possible to make good sense of relativism about truth—or so I hope to show.

In the first part of this paper, I will try to say exactly how the relativist's position should be stated. Relativism about truth, I will argue, is the view that truth (of sentences or propositions) is relative not just to contexts of use but also to *contexts of assessment*. The philosophical challenge is to explain what this talk of 'truth relative to a context of assessment' *means*. In the second part of the paper, I will meet this challenge—not by giving a definition (that is not how one should expect to illuminate fundamental concepts like truth), but by giving assessment-relative truth a role to play in a normative account of assertion. My account will settle precisely what one has to argue in order to defend the claim that a certain sentence or proposition is assessment-sensitive.

*Meeting of the Aristotelian Society, held in Senate House, University of London, on Monday, 6 June, 2005 at 4.15 p.m.

I

Stating the Position. One might think that being a relativist is just a matter of relativizing truth to some parameter. But it is not that simple. Many relativizations of truth are entirely orthodox. In model theory we talk of sentences being true relative to a model and an assignment of values to the variables, and in formal semantics we talk of sentences being true relative to a speaker and time, or more generally (following Kaplan 1989) a context of use. To my knowledge, no one has ever accused Tarski and Kaplan of being relativists for making use of these relativised forms of truth!

Sometimes relativism is presented as a thesis about sentence *tokens*: particular inscriptions or acoustic blasts. But even a sentence token can have different truth values on different occasions of use. When I leave my office for a quick errand, I put an old yellow post-it note with a token of 'I'll be back in a minute' on my door. Usually, this sentence token expresses a truth, but sometimes I get sidetracked and it expresses a falsehood.

Here the relativist might appeal to a distinction between sentence tokens and *utterances*. An utterance (in the sense relevant here) is an *act*. If I use my post-it note to announce that I will be back in a minute, my act counts as one utterance; if I do the same thing the next day, that is another utterance, using the same sentence token as a vehicle. The relativist thesis might be put this way: one and the same utterance or assertion can be true, relative to *X*, and false, relative to *Y*. This sounds more like a controversial thesis.

But there is something a bit odd about calling utterances or assertions, in the 'act' sense, true or false at all. We characterize actions as correct or incorrect, but not as true or false. We say 'His aim was true', but not 'His *aiming* was true', and it sounds equally funny to say 'That speech act was true' or 'What he did in uttering that sentence was true.' This suggests that when we say 'His assertion was false' or 'That was a true utterance', we are using 'assertion' and 'utterance' to refer to what was asserted or uttered, not to the act of asserting or uttering.¹ Characterizing relativism as a thesis about the truth

1. Bar-Hillel 1973, p. 304.

of assertions or utterances in the ‘act’ sense looks like a category mistake.

Similar considerations apply to the identification of *beliefs* as the things whose truth is ‘relative’. ‘Belief’ is ambiguous in much the same way as ‘assertion’. It can be used to refer to a *state* of a subject (Joe’s *believing* that newts are a kind of reptile) or to what is believed (that newts are a kind of reptile). When we say ‘Joe’s belief is true’, we are talking about the content of his belief, not the belief-state. That is why we can paraphrase ‘Joe’s belief is true’ as ‘What Joe believes is true’, but not as ‘Joe is in a true state.’

All of this suggests that the relativist doctrine should be stated as a claim about the truth of the things that are believed and asserted: propositions. Accordingly, Max Kölbel has suggested that

A relativism is not tame, if it involves the claim that the truth of propositions (or contents) of some kind can be relative, i.e. has the form

(RP) For any x that is a proposition of a certain kind K , it is relative to P whether x is true.²

But by this criterion, just about everyone who uses propositions in formal semantics would count as a non-tame relativist. For it is standard practice to relativise proposition truth to a *circumstance of evaluation*: typically a possible world, but in some frameworks a world and a time, or even a world and a standard of precision.³ The proposition that dodos are extinct in 2004 is true in the actual world, but there are possible worlds relative to which the very same proposition is false. Surely that does not vindicate relativism in any interesting sense.

Here it may seem tempting to say: ‘A *real* relativist is someone who takes proposition truth to be relative to some *other* parameter, in addition to worlds and possibly times.’⁴ But

2. Kölbel 2002, p. 119.

3. For recent discussions of what is at stake in choosing between these alternatives, see Richard 2003, Salmon 2003, and King, forthcoming. The debate turns on the logic of propositional attitude reports and on the proper treatment of tenses and expressions like ‘strictly speaking’, not on issues concerning relativism.

4. For this kind of response, see Nozick 2001, pp. 19, 307 n. 7.

this kind of response would miss the point. It's not the *kind* of parameter that matters, but how it's treated. We could even relativise proposition truth to an *aesthetic standards* parameter without being relativists about truth in any interesting sense! Suppose we took the predicate 'beautiful' to express a property whose extension varies not just with time and world but also with an aesthetic standard. We would presumably then say that the proposition expressed by

(1) Helen was beautiful at the beginning of the Trojan War

has truth values only relative to a world *and* an aesthetic standard. Would this make us relativists? Not all by itself. It depends on how the aesthetic standards parameter is treated in the definition of (sentence) truth at a context. An aesthetic absolutist might treat it this way:

Aesthetic absolutism: *S* is true at a context of use *C* iff there is a proposition *p* such that

- (a) *S* expresses *p* at *C*, and
- (b) *p* is true at the world of *C* and the One True Aesthetic Standard.

According to aesthetic absolutism, whether (1) is true at a context of use is completely independent of the speaker's (or anyone else's) aesthetic standards. The truth of (1) is no more 'relative' than the truth of any other tensed claim.

Why would an aesthetic absolutist *bother* relativising propositional truth to an aesthetic standard? For the same kinds of reasons that led Kaplan to relativise propositional truth to times.⁵ She might have independent semantic reasons for taking qualifiers like 'on any standard...' to be propositional operators. Propositional operators need a parameter to shift: just as modal operators shift the world parameter and temporal operators shift the time parameter, so these putative operators would shift the aesthetic standards parameter. I am not advocating this treatment of 'on any standard...'; I am merely pointing out that there might be good reasons for introducing a parameter of propositional truth that gets set to a constant

5. See Kaplan 1989, p. 502–4.

value in the definition of truth at a context. (Indeed, that is precisely how actualists treat the world parameter.)

Alternatively, one might look to the context of use to fix a value for the aesthetic standards parameter:

Aesthetic contextualism: S is true at a context of use C iff there is a proposition p such that

- (a) S expresses p at C , and
- (b) p is true at the world of C and the aesthetic standards of the speaker at C .

This is a kind of relativism about beauty, perhaps, but not about truth. There is always an absolute answer to the question, 'Is S true at C ?' or 'Did A utter S truly?'

How would one have to treat the aesthetic standards parameter in order to be a relativist about truth? Here is my suggestion: the parameter would have to be initialised not by a constant (as in aesthetic absolutism) or by a feature of the context of use (as in aesthetic contextualism), but by a feature of the context in which the speech act (or other use of the sentence) is being *assessed*. In order to state the relativist's position, then, we must employ the doubly contextual predicate 'true at context of use C_U and context of assessment C_A ' in place of the familiar 'true at context of use C '. By a 'context of assessment', I mean simply a concrete situation in which a use of the sentence is being *assessed*. We perform speech acts, but we also assess them; so, just as we can talk of the context in which a sentence is being used, we can talk of a context (there will be indefinitely many) in which a use of it is being assessed. Using this notion, we can articulate a radical aesthetic relativism:

Aesthetic relativism: S is true at a context of use C_U and context of assessment C_A iff there is a proposition p such that

- (a) S expresses p at C_U , and
- (b) p is true at the world of C_U and the aesthetic standards of the assessor at C_A .

This formulation does, I think, capture what the relativist is after. Whether we can make any sense of the doubly contextual truth

predicate to which it appeals is another question, to which we will soon turn.

First, however, it will be useful to fix some terminology. I will call a sentence *use-sensitive* if its truth value varies with the context of use (keeping the context of assessment fixed), *assessment-sensitive* if its truth value varies with the context of assessment (keeping the context of use fixed), and *context-sensitive* if it is either use-sensitive or assessment-sensitive. Similarly, I will call a sentence *use-indexical* if it expresses different propositions at different contexts of use (keeping the context of assessment fixed), *assessment-indexical* if it expresses different propositions at different contexts of assessment (keeping the context of use fixed), and *indexical* if it is either use-indexical or assessment-indexical.

Although indexicality and context sensitivity are often conflated, it is important to distinguish them. A sentence can be use-sensitive without being use-indexical: consider, for example,

- (2) The number of AIDS-infected babies born in Oakland in 2004 is 65.

Although (2) is use-sensitive—its truth at a context of use depends on the *world* of the context—it is not use-indexical: it expresses the same proposition at every context of use.⁶ Use indexicality and use sensitivity can come apart because the context of use plays two distinct roles in determining the truth value of a sentence at a context, as can be seen from Kaplan's definition of sentence truth at a context:

If *c* is a context, then an occurrence of [a sentence] ϕ in *c* is true iff the content expressed by ϕ in this context is true when evaluated with respect to the circumstance of the context.⁷

First, the context of use helps determine which proposition is expressed. But because this proposition has truth values only relative to circumstances of evaluation, we must appeal to the

6. Of course there are worlds in which sentences orthographically identical with (2) have completely different meanings, or no meanings at all. But when we ask what proposition a sentence expresses relative to a context of use, we are asking about the sentence with its actual meaning.

7. Kaplan 1989, p. 522; for a formal version, see p. 547.

context of use a second time to fix the relevant circumstance of evaluation (what Kaplan calls ‘the circumstance of the context’). In a sentence that is use-sensitive but not use-indexical, context plays no role at the first step but still has an effect on truth value at the second.⁸

For the same reason, assessment sensitivity need not be due to assessment indexicality. According to the version of aesthetic relativism presented above, for example, (1) is assessment-sensitive but not assessment-indexical. We do not need to allow that the propositional content of an assertion might vary from one assessor to another in order to make sense of relative truth.

The concepts of use sensitivity and assessment sensitivity can be applied to propositions as well as sentences. If it seems odd to characterize a *proposition* as context-sensitive, remember that not all context sensitivity is due to indexicality. (2) is use-sensitive not because it expresses different propositions at different contexts of use, but because the proposition it expresses is itself use-sensitive: this proposition can be truly asserted at some contexts but not at others. I will say that a *proposition* is true at a context of use C_U and context of assessment C_A just in case it is true at the circumstance of evaluation determined by C_U and C_A .⁹ (The relation of ‘determination’ will vary from one semantic theory to another: in the aesthetic relativist’s semantics, for example, the circumstance of evaluation determined by C_U and C_A will be composed of the world of C_U and the aesthetic standards of the assessor at C_A , while in the aesthetic contextualist’s semantics, it will be composed of the world of C_U and the aesthetic standards of the speaker at C_U .) I will call a proposition *use-sensitive* if its truth value varies with the context of use (keeping the context of assessment fixed), and *assessment-sensitive* if its truth value

8. Distinguishing these two roles for context helps illuminate what is at stake in the debate between ‘eternalists’ (like Salmon and Richard) and ‘temporalists’ (like Kaplan). The issue is whether the time of utterance affects the truth values of tensed sentences in the first way (by determining which proposition is expressed) or the second (by determining how it is to be evaluated).

9. This assumes that a context of use and context of assessment will always determine a unique circumstance of evaluation. There are some semantic applications for which that restriction is too limiting. A more general formulation would replace ‘the circumstances of evaluation determined by C_U and C_A ’ with ‘all the circumstances of evaluation compatible with C_U and C_A ’.

varies with the context of assessment (keeping the context of use fixed).

We are now in a position to state the relativist's position in its full generality. *Relativism about truth* is the view that there is at least one assessment-sensitive sentence. If we restrict the domain to natural languages, or to some particular language, we get a thesis that is at least partly empirical, while if we broaden it to all conceivable languages, we get a thesis that might be settled *a priori*. Two further subdivisions are useful. An *expressive relativist* holds that there is at least one assessment-indexical sentence, while a *propositional relativist* holds that there is at least one assessment-sensitive proposition. In what follows, I will focus on propositional relativism, which seems to me more promising in its applications.

II

What Does It Mean? We have stated the relativist's thesis. But do we really understand it? We cannot understand it unless we grasp what is meant by 'true at context of use C_U and context of assessment C_A ', and it is not clear that we do. For it is not clear that the concept of truth *admits* of relativisation to assessors. If 'true' as it occurs in 'true for X ' is just the ordinary, non-relative truth predicate, then it is unclear what 'for X ' adds, unless it is just 'and X believes this'. On the other hand, if the occurrence of 'true' in 'true for X ' is like the 'cat' in 'cattle', then the relativist needs to explain what 'true for X ' means and what it has to do with truth, as ordinarily conceived.¹⁰

Relativists often try to meet this challenge by giving a *definition* of truth that makes its relativity plain. If truth is idealized justification, then it might reasonably be thought to be assessor-relative, since ideal reasoners with different starting beliefs or prior probabilities might take the same ideal body of evidence to support different conclusions. Similarly, if truth is defined pragmatically, as what is good to believe, then it might also be assessor-relative, insofar as different things are good for

10. See Meiland 1977. Meiland's own explication of 'true for X ' as 'corresponds to reality for X ' just pushes the problem back a level. The absolutist can object that her understanding of 'correspondence to reality' leaves no room for an added 'for X '.

different assessors to believe. But although these coherentist and pragmatic definitions of truth capture the ‘relative’ part of ‘relative truth’, I do not believe they capture the ‘truth’ part. Indeed, for familiar reasons, I doubt that the concept of truth can be usefully illuminated by a definition in terms of more primitive concepts.¹¹

Of course, the relativist semanticist can give a formal definition of ‘true at context of use C_U and context of assessment C_A ’ that fixes its extension over a particular class of sentences and contexts. But such a definition would not answer the challenge, for reasons Michael Dummett made clear in his classic paper ‘Truth’ (1959). Dummett pointed out that a set of T-biconditionals or a recursive definition of ‘true in L ’ cannot simultaneously explain the meanings of the expressions of L and the meaning of ‘true in L ’. Thus, if our aim is to explain the meanings of expressions by showing how they contribute to the truth conditions of sentences containing them, we must have a grasp of the concept of truth that goes beyond what a Tarskian truth definition tells us. On Dummett’s view, this grasp consists (at least in part) in our knowledge that the central convention governing the speech act of assertion is to assert only what is true.

Two things are worth noting here. First, if Dummett is right, then it is not just the relativist who owes an explication of the significance of her truth predicate. The absolutist owes one as well—at least if she is to use this predicate in semantics. Second, Dummett’s proposed explication does not take the form of a definition. Instead, it is a description of the role ‘true’ plays in a broader theory of language use: specifically, an account of the speech act of assertion. These two points suggest a strategy for the relativist: start with such an explication of truth (one that is acceptable to the non-relativist), then find a job for contexts of assessment in this framework. Having done this, the relativist should be able to say to the absolutist: ‘If you can make sense of your absolute truth predicate, you should be able to make sense of my relative one, too, and to see why it deserves to be called a *truth* predicate.’

11. See Davidson 1997.

The strategy is, I think, a promising one. But if we try to execute it by generalizing Dummett's explication of truth as the conventional aim of assertion, we immediately run into difficulties—difficulties that may explain why so many philosophers have dismissed relative truth as unintelligible. There are three ways in which we might connect the aim of assertion with doubly contextual truth:

1. *Relativise the aim of assertion to contexts of assessment:* Relative to context C_A , assertion is governed by the convention that one should assert at context C_U only what is true relative to context of use C_U and context of assessment C_A .
2. *Quantify over contexts:* One should assert at C_U only what is true at context of use C_U and *some/most/all* contexts of assessment C_A .
3. *Privilege one context of assessment* (the one occupied by the asserter at the moment of utterance): One should assert at C_U only what is true at context of use C_U and context of assessment C_U .

But none of these options will help us understand assessment-relative truth.

Option 1 just replaces one unexplicated relativisation with another. Conventions supervene on patterns of mutual belief and expectation among the participants in a practice.¹² So the only way conventions can be assessment-relative is if facts about the participants' mental states are assessment-relative. An explication of assessment-relative truth that presupposes an understanding of assessment-relative *facts* is not going to get us very far.

Option 2 is at least intelligible, but it does not serve the relativist's purposes. It is too easy to assert something that is true at *some* context of assessment, and if we require truth at *every* context of assessment, the resulting norm will forbid asserting *anything* assessment-sensitive. 'Most' seems the best choice of quantifier, but there is something arbitrary about it; majority rule looks misplaced here. Nor is it clear what 'most' means in this context, if, as seems likely, there are infinitely many possible

12. See Lewis 1969.

contexts of assessment. More seriously, option 2 leaves room for the anti-relativist to respond by saying:

What you call ‘truth at some/all/most contexts of assessment’ and identify with the aim of assertion is what I call ‘truth (*simpliciter*)’. But I fail to see what you mean by ‘true at context of assessment C ’. Suppose there are only three possible contexts of assessment: C_1 , C_2 , and C_3 . What is the practical difference between being true at C_1 and C_2 but not C_3 and being true at C_2 and C_3 but not C_1 ? Nothing you have said discriminates between these possibilities, so I still lack any understanding of the difference.

Option 3 is the choice of most relativists who have considered the matter at all.¹³ If a single context of assessment is to be privileged as that relative to which one should assert only truths, it seems reasonable that it should be the context one occupies when making the assertion. But a version of the previous objection applies here as well. Option 3 gives a significance to ‘true at context of use C_U and context of assessment C_A ’ only for the special case where $C_U = C_A$, and not for arbitrary C_U and C_A . As a result, it cannot help us to understand assessment sensitivity. Suppose two rival semantic theories, T_1 and T_2 , agree about the truth value of S at a context of use C_U and context of assessment C_A whenever $C_U = C_A$, but disagree about the truth value of S at C_U and C_A for at least some context pairs such that $C_U \neq C_A$. According to T_1 , S is assessment-sensitive, while according to T_2 , it is not (see Fig. 1).

The relativist ought to have something to say about how these theories differ from each other in practice and how one might decide between them. But if all we’re told is that the aim of assertion is to assert something that is true as assessed from the context of use, we cannot discern any practical difference between T_1 and T_2 .

It might be protested that although the difference between T_1 and T_2 does not manifest itself as a difference in the norms for asserting S , it manifests itself as a difference in the norms for asserting that particular utterances of S are ‘true’. For example, T_1 and T_2 disagree about whether it would be correct

13. See Kölbel 2002, p. 125 and Egan, Hawthorne, and Weatherson, forthcoming.

T_1	$C_A=$	T_2	$C_A=$
	C_1 C_2 C_3		C_1 C_2 C_3
C_1	T	T	F
$C_U=$ C_2	T	T	F
C_3	T	F	F

C_1	T	T	T
$C_U=$ C_2	T	T	T
C_3	F	F	F

Figure 1: T_1 and T_2 .

for a speaker at context C_1 to call an utterance of S at C_3 ‘true’. But this reply puts the cart before the horse. It is important to distinguish the monadic predicate ‘true’, which is just another word in the language being studied, from the three-place predicate ‘true at context of use ... and context of assessment ...’, which the semanticist uses in describing the language. It would be mad to explain the main semantic predicate in our meta-language by appealing to the use of an *object-language* expression whose meaning we are describing using that very predicate. And what happens if the language we are studying does not contain ‘true’ or its equivalent? Do we then lose our grip on the significance of relative-truth ascriptions?¹⁴

I think we must conclude that there is no prospect of generalizing the Dummettian conception of truth as the aim of assertion in a way that makes assessment-relative truth intelligible. But it would be too hasty to conclude from this that relative truth talk is incoherent. We have only explored one approach to explicating truth talk (Dummett’s), and we might reject this approach for reasons that have nothing to do with assessment sensitivity. After all, truth is not the *only* thing we are conventionally understood to be aiming at in making assertions. We also expect assertions to be warranted by what the asserter knows and relevant to the conversational setting in which they

14. See Dummett 1981, pp. 320–1. The issues here are subtle, in part because in a language that contains any assessment-sensitive expressions, the predicate ‘true’ must also be assessment-sensitive.

occur. Moreover, we expect asserters to be sincere: to assert only what they believe. So truth can hardly be singled out as *the* conventional aim of assertion. It is not even obvious that it is *a* conventional aim of assertion. An insincere assertion that happens to be true seems a more flagrant violation of the norms for assertion than a sincere (and warranted) one that happens to be false. Perhaps assertion aims at truth only indirectly, by aiming at the sincere expression of *belief*, which aims in turn at truth.

The claim that belief ‘constitutively’ aims at truth is on firmer ground: it is arguable that a cognitive state that did not aim at truth would not count as a belief at all.¹⁵ We might accordingly try to understand truth as the condition for the correctness of beliefs. The problem is that (as with assertions) there are many dimensions along which beliefs might be assessed as correct or incorrect. If you have patiently gathered the evidence and it overwhelmingly favours not-*p*, there is an important sense in which it would be incorrect for you to believe that *p*, even if *p* happens to be true. But there is also a sense in which your belief that *p* would be correct. One might try to distinguish the two senses of correctness at issue here—perhaps as ‘subjective’ and ‘objective’, or as ‘epistemic’ and ‘representational’—but it is hard to see how this could be done without invoking the notion of truth. It seems unlikely, then, that one could get a grip on the concept of truth just by being told that truth is the aim of belief, or the norm for correct belief.

How, then, *should* we understand the significance of truth talk in semantic theorizing? I think Dummett is right that our grip on truth comes from an understanding of its relation to assertion. But where Dummett focused on the norms *for* making an assertion, I propose we focus on the normative *consequences* of making an assertion. An assertion (even an insincere one) is a *commitment to the truth* of the proposition asserted.¹⁶ It might be thought that this idea is just as inimical to relative truth as Dummett’s. Just as it doesn’t make sense to aim to speak truth if truth is relative, so (one might suppose) it doesn’t make sense to commit oneself to the truth of a proposition if truth

15. See Williams 1973.

16. See for example Searle 1979, p. 12, though the idea is ubiquitous.

is relative.¹⁷ But that's not clear—partly because it's not clear what it *means* to be committed to the truth of a proposition. How, exactly, does one honour or violate such a commitment? What is one committed to *doing*? I want to suggest that when this is spelled out in a plausible way, we can make very good sense of 'commitment to truth' even if truth is assessment-relative.

Here are three things that might be thought to constitute the 'commitment to truth' one undertakes in making an assertion:

- (W) Commitment to withdraw the assertion if and when it is shown to have been untrue.
- (J) Commitment to justify the assertion (provide grounds for its truth) if and when it is appropriately challenged.
- (R) Commitment to be held responsible if someone else acts on or reasons from what is asserted, and it proves to have been untrue.

Everyone should be able to agree that assertoric commitment includes at least (W). Imagine someone saying: 'I concede that what I asserted wasn't true, but I stand by what I said anyway.' We would have a very difficult time taking such a person seriously as an asserter. If she continued to manifest this kind of indifference to established truth, we would stop regarding the noises coming out of her mouth as assertions. We might continue to regard them as expressions of beliefs and other attitudes (just as we might regard a dog's whining as an expression of a desire for food). We might even find them useful sources of information. But we would not regard them as commitments to truth, and hence not as assertions.

There will be less agreement about (J). Brandom has argued that assertoric commitment includes (J) as well as (W) (Brandom 1983, 1994), but this may be over-generalizing from seminar-room assertions to assertions in general. Suppose someone were to say: 'You've given some very good reasons to doubt the truth of what I asserted. I have nothing to say in answer to your objections, yet I continue to stand by my claim.' She would not be playing the game of assertion the way philosophers play it,

17. Burnyeat argues that the Protagorean relativist is thwarted by 'the commitment to truth absolute which is bound up with the very act of assertion' (1976, p. 195).

but perhaps philosophers do not get to set the rules here. We would surely take her assertions less seriously than we would if she were responsive to reasons. But would we cease treating her as an asserter at all? That is not so clear.

What about (R)? Asserting is a bit like *giving one's word* that something is so, and our reactions to assertions that turn out to have been untrue can resemble our reactions to broken promises. We feel a legitimate sense of grievance, especially if we have acted on what we were told. Suppose someone tells you that there will be a talk by an interesting celebrity at a nearby university. You cancel some appointments and spend considerable time and energy getting there—but there is no talk. You expect that when you confront your informant, she will apologize profusely. Even if she has an excuse (perhaps there was a typo in the schedule), she will accept some measure of responsibility. You will be shocked if she says: 'You actually *acted* on my assertion? Well, that's not my problem. It's up to *you* to sort out what's worth taking seriously.' But why does this response sound so wrong? After all, it *is* up to us whether to believe what we are told, and we don't expect our informants to be infallible. A plausible answer (though not the only one) is that part of what it is to make an assertion is to accept partial responsibility for the accuracy of what one says.

Suppose we understand the assertoric 'commitment to truth' in terms of some combination of (W), (J), and (R). Can we understand what it would be to commit oneself to the truth of an assessment-sensitive proposition? That is, can we find plausible construals of (W), (J), and (R) in a framework that relativises truth to contexts of assessment as well as contexts of use?

(W) talks of the asserted proposition being 'shown to have been untrue'. Untrue, relative to which context of use and context of assessment? The relevant context of use is obviously the context in which the proposition was asserted. But what about the context of assessment? There are four natural options:

1. Quantify over contexts of assessment: the proposition must be shown to be untrue relative to the context of use and *some/all/most* contexts of assessment.
2. The relevant context of assessment is the context in which the proposition was asserted (= the context of use).

3. The relevant context of assessment is the context in which the putative refutation is being given.
4. The relevant context of assessment is the context in which the asserter is evaluating the putative refutation.

It should be plain from our parallel discussion of the aim of assertion that the first two options will not help make sense of assessment sensitivity. They imply that for any given assessment-sensitive proposition, there will be a systematically related assessment-invariant proposition whose assertion results in exactly the same commitments. But unless we can see some difference in practice between asserting an assessment-sensitive proposition and asserting a related assessment-invariant one, we lack a real understanding of assessment sensitivity.

Only the third and fourth options give an essential and ineliminable role to contexts of assessment. They differ on what a successful refutation must establish: while the fourth option demands proven untruth relative to the *asserter's* context, the third looks to the *challenger's* context. The third option can be ruled out, I think, as too damaging to the integrity of a single person's body of assertions. If I withdraw some of my assertions because they are untrue relative to Bob's context and others because they are untrue relative to Marie's, I may end up with a body of assertions that is incoherent and reflects no one's point of view. This would be a bit like letting a bush be pruned by several gardeners with radically different conceptions of how it should look: the little that remained would not satisfy any of them. It demands too much of asserters to give every challenger the home stadium advantage.

I conclude that the relativist should construe (W) along the lines of the fourth option, which privileges contexts the asserter occupies, while still allowing the relevant context of assessment to diverge from the context of use:

- (W*) In asserting that p at C_1 , one commits oneself to withdrawing the assertion (in any future context C_2) if p is shown to be untrue relative to context of use C_1 and context of assessment C_2 .

There should be no worries about the intelligibility of (W*). Logically, it is no more complex than a commitment to refill

the pitcher (at any future time t) if it is shown to be empty (at t). And it reduces to the original (W) when p is not assessment-sensitive.

It should now be clear how we must generalize (J). Since (W*) requires that an assertion be withdrawn when it is proven untrue relative to the asserter's current context of assessment, the justification demanded by (J*) must consist in grounds for the truth of the asserted proposition relative to this same context:

- (J*) In asserting that p at C_1 , one commits oneself to justifying the assertion when the assertion is appropriately challenged. To justify the assertion in a context C_2 is to provide grounds for the truth of p relative to context of use C_1 and context of assessment C_2 .

For similar reasons, we must construe (R) as follows:

- (R*) In asserting that p at C_1 , one commits oneself to accepting responsibility (at any future context C_2) if on the basis of this assertion someone else takes p to be true (relative to context of use C_1 and context of assessment C_2) and it proves to be untrue (relative to C_1 and C_2).

Let's take stock. We began with the worry that we did not really *understand* the relativist's doubly contextual truth predicate, 'true at context of use C_U and context of assessment C_A '. To assuage this worry, we decided, it would not be necessary to give an informative definition of this predicate in conceptually simpler terms (since not even the absolutist can do *that*): it would be enough to describe the predicate's role in a larger theory of meaning. We have now done just this. We have given an account of assertoric commitment that settles just what one is committing oneself to in asserting an assessment-sensitive proposition. By doing this, I suggest, we have made relativism about truth intelligible.¹⁸

18. Once we have come to understand the relativist's doubly contextual truth predicate through reflection on its connection with assertion, we can employ it in our theories of the propositional attitudes and speech acts other than assertion without worrying that our use of it is completely unconstrained.

But is relativism *true*? The weakest relativist position we distinguished in Section I was the claim that there is at least one assessment-sensitive sentence in some conceivable language. We have already said enough to vindicate this claim. We can certainly imagine a language in which the word ‘beautiful’ works as described by the aesthetic relativist (from Section I): just imagine that its speakers use sentences containing ‘beautiful’ to undertake the commitments implied by (W*), (J*), and (R*), together with the relativist’s semantics. (Some philosophers may hold that English is such a language.) Even if you think it would be rash or irresponsible to undertake such open-ended commitments, it seems at least conceivable that speakers might do so, and that they might have a conventional linguistic way of doing so. So the weakest form of relativism about truth would seem to be true.¹⁹

A stronger, more interesting thesis is that some of the things *we* say and think are assessment-sensitive. We have not established that, but we have at least shown what such a claim would imply and what evidence might count for or against it. To defend an assessment-sensitive semantics for a particular class of sentences, one would have to adduce evidence about the norms for defending and withdrawing assertions made using those sentences. I think that quite a good case can be made for the assessment sensitivity of the future tense, epistemic modals, knowledge attributions, predicates of personal taste, and other constructions, but I cannot argue that here.^{20,21}

Department of Philosophy
University of California
Berkeley, CA 94720-2390
USA
jgm@berkeley.edu

19. There are no worries about self-refutation here, because we may suppose ourselves to be describing the ‘relativistic’ language in a meta-language devoid of assessment sensitivity.

20. For some arguments, see Kölbel 2002, 2004; MacFarlane 2003, forthcoming; Richard 2004; and Egan, Hawthorne, and Weatherson, forthcoming.

21. I am grateful to audiences at Harvard, Princeton, Santa Barbara, and San Diego for useful feedback on ancestors of this paper. This research was supported in part by an ACLS/Andrew W. Mellon Fellowship for Junior Faculty and a UC Berkeley Humanities Research Fellowship.

REFERENCES

- Bar-Hillel, Y. 1973: 'Primary Truth Bearers', *Dialectica* 27, pp. 303–12.
- Brandom, R. 1983: 'Asserting', *Noûs* 17, pp. 637–50.
- Brandom, R. 1994: *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Cambridge, Massachusetts: Harvard University Press.
- Burnyeat, M. F. 1976: 'Protagoras and Self-Refutation in Plato's *Theaetetus*', *Philosophical Review* 85, pp. 172–95.
- Davidson, D. 1997: 'The Folly of Trying to Define Truth', *Journal of Philosophy* 94, pp. 263–78.
- Dummett, M. 1959: 'Truth', *Proceedings of the Aristotelian Society*, n.s. 59, pp. 141–62.
- Dummett, M. 1981: *Frege: Philosophy of Language*, 2nd edition, Cambridge, Massachusetts: Harvard University Press.
- Egan, A., J. Hawthorne, and B. Weatherson forthcoming: 'Epistemic Modals in Context', in G. Preyer and P. Peter (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press.
- Kaplan, D. 1989: '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.
- King, J. forthcoming: 'Tense, Modality, and Semantic Value', *Philosophical Perspectives*.
- Kölbel, M. 2002: *Truth Without Objectivity*, London: Routledge.
- Kölbel, M. 2004: 'Faultless Disagreement', *Proceedings of the Aristotelian Society* 104, pp. 53–73.
- Lewis, D. 1969: *Convention*, Cambridge, Massachusetts: Harvard University Press.
- MacFarlane, J. 2003: 'Future Contingents and Relative Truth', *Philosophical Quarterly* 53, pp. 321–36.
- MacFarlane, J. forthcoming: 'The Assessment Sensitivity of Knowledge Attributions', *Oxford Studies in Epistemology*.
- Meiland, J. 1977: 'Concepts of Relative Truth', *The Monist* 60, pp. 568–82.
- Nozick, R. 2001: *Invariances: The Structure of the Objective World*, Cambridge, Massachusetts: Harvard University Press.
- Richard, M. 2003: 'Introduction to Part I', in A. Jokić and Q. Smith (eds.), *Time, Tense, and Reference*, Cambridge, Massachusetts: MIT Press.
- Richard, M. 2004: 'Contextualism and Relativism', *Philosophical Studies* 119, pp. 215–42.
- Salmon, N. 2003: 'Tense and Intension', in A. Jokić and Q. Smith (eds.), *Time, Tense, and Reference*, Cambridge, Massachusetts: MIT Press.
- Searle, J. R. 1979: *Expression and Meaning*, Cambridge: Cambridge University Press.
- Williams, B. 1973: 'Deciding to Believe', in *Problems of the Self*, Cambridge: Cambridge University Press.