

Relativism and Disagreement

John MacFarlane

August 17, 2006

It has often been proposed that claims about what is funny, delicious, or likely are “subjective,” in the sense that their truth depends not only on how things are with the objects they explicitly concern, but on how things are with some *subject* not explicitly mentioned. This thought is supported by the striking degree to which we differ in our judgements about these matters. If there are wholly objective properties of funniness, deliciousness, or likelihood, then most of us must be *defective* in our capacity to detect them. We are humor-blind, or taste-blind, or likelihood-blind, in much the same way that some of us are color-blind. But this diagnosis clashes with the way we think and talk about these domains. In our judgements about what is delicious, we lack the humility color-blind people show in their judgements about what is red or green. We do not seem to regard the fact that many others disagree with us as grounds for caution in calling foods delicious. We readily judge things to be funny in light of our own senses of humor, even though when challenged we can offer no grounds for thinking our senses of humor are the “right” ones. We readily judge things to be likely in light of what we know, even while acknowledging that our knowledge is only partial, and that others may know more than we do.

There is more to say to the hard-core objectivist, but for purposes of this paper let us dismiss her here. Let us also dismiss the philosopher who denies that claims about what is funny, delicious, or likely are in the right line of work to be counted true or false.

Here we will be concerned with two competing ways in which the idea that a discourse is “subjective” might be cashed out semantically: contextualism and relativism.

The relativist’s central objection to contextualism is that it fails to account for the possibility of *disagreement* in subjective discourse—for our sense that when I say that carrots are delicious and you deny this, we are genuinely *disagreeing* with each other, and not making compatible claims about our respective tastes. If we are to adjudicate between contextualism and relativism, then, we must first get clear about what it is for two people to disagree. This question turns out to be surprisingly difficult to answer. Although the answer given below will be incomplete, I think it does shed a little light on what the relativist must say if she is to do better than the contextualist in securing genuine disagreement.

1 Disagreement Lost

1.1 Contextualism

The contextualist takes the subjectivity of a discourse to consist in the fact that it is covertly about the speaker (or perhaps a larger group picked out by the speaker’s context and intentions). Thus, in saying that apples are “delicious,” the speaker says, in effect, that apples taste good to her (or to those in her group). In saying that a joke is “funny,” she says that it appropriately engages her sense of humor (or that of her group). And in saying that some state of affairs is “likely,” she says that it is likely given what she knows, or perhaps what she and her co-investigators know.

This kind of view has obvious appeal. It explains how the truth of the claims at issue can depend both on how things are with their explicit subject-matter (say, apples) and on how things are with a subject or subjects who is not explicitly mentioned. And it does

so in a perfectly straightforward way, invoking semantic mechanisms that are already needed to handle the more familiar kinds of context sensitivity exhibited by indexicals and demonstratives, quantifiers, and gradable adjectives.

1.2 The problem of lost disagreement

But the contextualist solution has a price. If in saying “apples are delicious” I am saying that they taste good *to me*, while in saying “apples are not delicious” you are denying that they taste good *to you*, then we are no more disagreeing with each other than we would be if I were to say “My name is John” and you were to say “My name is not John.” Intuitively, though, it does seem that we are disagreeing. We certainly *take ourselves* to be disagreeing. I may say, “Wrong!” or “That’s false”—neither of which would be appropriate if you had said explicitly that apples taste good to you.

Similarly, if I say “It’s likely that the thieves are in Mexico by now,” you may reply, “No, it’s not likely at all: I just heard that they were seen heading north towards Oregon.” You may concede that, given the information I had when I spoke, it was probable that the thieves were in Mexico. Nonetheless, it seems quite natural for you to contradict my claim in light of information only you possess. It also seems natural for me to retract my claim when you do. None of this makes much sense on the contextualist analysis, which takes my claim to be about what was possible in light of the information I had at the time I made it. (Of course, there are moves the contextualist can make here. We will consider them shortly.)

Thus the contextualist saves some of the appearances, but not all. The contextualist can explain, as the hard-core objectivist cannot, why speakers so readily make claims in these domains on the basis of their own idiosyncratic tastes, senses of humor, or knowledge. However, by construing these claims as claims about the speaker (or some contextu-

ally relevant group), the contextualist makes it difficult to make sense of the *disagreement* speakers perceive in these areas of discourse.

1.3 Single scoreboards: sharing the subjectivity

Even contextualists acknowledge the need to account for intuitions of disagreement. Thus, Keith DeRose writes (concerning knowledge attributions):

...the skeptic and her opponent do *take themselves* to be contradicting one another; each *intends* to be contradicting what the other is saying; and, beyond what's going on privately in their own minds, each is publicly indicating that they are (or at least mean to be) contradicting the other, by saying such things as, 'No, you're wrong. I *do* know.' ... And about this kind of case, I am *not* inclined to think that both our speakers are speaking the truth, but failing to contradict one another. (DeRose 2004, 3)

So how can a contextualist do justice to our sense that the parties to the conversations we have described are disagreeing with each other? DeRose's approach is to invoke David Lewis's (1979) metaphor of a "scoreboard" that changes in response to various conversational moves. Instead of taking partially subjective discourse to be (partly and covertly) about the speaker, we take it to be (partly and covertly) about some aspect of the shared conversational score. Thus, for example, the scoreboard might include a shared epistemic standard that changes as the conversation evolves, getting more stringent when the stakes are high and less stringent when they are low. The contextualist about knowledge attributions can then take an occurrence of "*S* knows that *p*" at a context *C* to be true only if *S* stands in a good enough epistemic position with respect to *p* to satisfy the epistemic standard on the shared conversational scoreboard at *C*. This kind of contextualism avoids the problem of lost disagreement that would be threatened if "*S*

knows that p ” always expressed a claim about the speaker’s personal epistemic standards. As DeRose puts it: “If there is a single scoreboard, then our skeptic is denying precisely what her opponent is affirming as they debate back and forth” (2004, 6).

Parallel moves are possible in other domains. For example, claims about what is “funny” can be construed as claims about what tickles the sense of humor that is taken to be shared by the conversational participants (at least for purposes of conversation—they need not actually share it). And claims about what is “likely” can be construed as claims about what is probable given what is known by the coinvestigators.¹ Working out such views in detail would require specifying how the shared conversational score is determined by the mental states and linguistic performances of the participants, and what happens when there is not enough agreement to establish a single shared score. DeRose has an interesting discussion of the options, but these details won’t matter for our purposes.

The single-scoreboard gambit attempts to strike a compromise between two competing desiderata. On the one hand, we want to capture the *subjectivity* of the claims at issue—the degree to which their truth seems to depend not just on how things are with the objects they are explicitly about, but on how things are with certain subjects. On the other hand, we want to capture the relations of *disagreement* that speakers perceive between different claims in these domains. The hard-core objectivist captures the disagreement perfectly, but loses the subjectivity. The simple contextualist captures the subjectivity perfectly, but loses the disagreement. The single-scoreboard approach tries to strike a balance between these desiderata by taking the subjectivity to be *shared* among the participants of a single conversation.

On the subjectivity side, this compromise seems acceptable. In many cases, we do pay attention to “shared score” in deciding whether to assert that something is “funny” or

“delicious” or “likely.” There are contexts in which speakers will refrain from calling a joke “funny” but still be prepared to assert, “*I* was really amused by it.” Similarly, there are contexts in which, although it is probable given what *I* know that *p*, I may refrain from saying that *p* is “likely,” because I suspect that one of the other group members may have knowledge that counts against *p*.

Granted, there are also cases in which two parties will continue to disagree over whether something is “delicious” even after it has become clear that their tastes are not consonant enough to establish a single shared standard of taste for the conversation. One might wonder what they could possibly be arguing about, if (as the single scoreboard view has it) their claims concern a shared standard of taste. Why don’t they give up once they realize that there *is* no shared standard? But it seems to me that the single scoreboarder has a decent response here: even if there is no hope of either assertion being *true* (given that there is no established common standard of taste), the speakers might make these assertions as a way of trying to *establish* a common standard of taste, through what Lewis (1979) has called “accommodation.”

So the single-scoreboard compromise does pretty well with subjectivity. The problem is that it doesn’t give us enough disagreement. It gives us disagreement only within the bounds of a single “conversation”—something for which it makes sense to imagine a shared scoreboard. But disagreement isn’t confined to these bounds. Consider these examples:

1. When I was ten, I used to go around saying “fish sticks are delicious” (and meaning it!). Now I say “fish sticks are not delicious.” It seems to me that I disagree with my past self. But I am not involved in a *conversation* with my past self.
2. Similarly, when I overhear a group of ten-year-olds chattering about how “funny” certain knock-knock jokes are, I may think that they are wrong. These jokes just

aren't that funny. But the kids certainly don't think of themselves as involved in a conversation with *me*—they may not even know I'm there. Nor do I think of myself as conversing with them.

3. Unbeknownst to each other, two research groups are investigating the cause of a certain epidemic. After gathering lots of evidence, Group 1 concludes, "It's likely that the epidemic was spread to humans from birds." Group 2, looking at a different body of evidence, concludes, "It's not likely that the epidemic was spread from birds; it probably came from spores in the soil." Looking at both reports, I note that their conclusions *disagree* about the likelihood that the disease was spread to humans from birds. When I bring both groups together, they will accept my characterization of their results as disagreeing. They will certainly not protest that their conclusions are compatible.

The problem with the single scoreboard approach is that it explains only *intra-conversational* disagreement, leaving inter-conversational disagreement unaccounted for. This is not a stable resting point. Once the importance of accounting for disagreement has been conceded, one cannot limit oneself to disagreement within conversations.² And it is hopeless to widen the bounds of "conversations" as needed to make *all* disagreement intra-conversational. For it is only if conversations are bounded and relatively self-contained that we can really make sense of the idea of a shared scoreboard.

2 Disagreement Regained?

One might now despair of ever getting subjectivity and disagreement into the same picture. Perhaps we just have to choose. This is where the relativist comes in with her seductive song. "You can have it both ways," she says, "if you just accept that propo-

sitions about the funny, the delicious, and the likely have truth values only relative to a person or perspective. When I say that apples are delicious and you deny this, you are denying the very same proposition that I am asserting. We genuinely disagree. Yet this proposition may be true for you and false for me. That is what the ‘subjectivity’ of these claims comes to: perspectival dependence of their truth on features of the subjects who assess them.”

Such a view might not be as wild as it sounds. On standard views, propositions have truth values relative to “circumstances of evaluation.” These are typically taken to be possible worlds, but a minority tradition takes circumstances of evaluation to be world/time pairs. And Kaplan explicitly leaves room for other parameters:

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. (1989, 502)

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 neutral. 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)

Taking this line of thought a little farther, the relativist might envision contents that are “sense-of-humor neutral” or “standard-of-taste neutral” or “epistemic-state neutral,” and circumstances of evaluation that include parameters for a sense of humor, a standard of taste, or an epistemic state. This move would open up room for the truth value of a proposition to vary with these “subjective” factors in much the same way that it varies

with the world of evaluation. The very same proposition—say, that apples are delicious—could be true with respect to one standard of taste, false with respect to another.

The question I want to address is whether this helps with the problem of lost disagreement. How can there be *disagreement* between Abe and Ben, on the relativist's view, if the proposition Abe asserts and Ben denies is true relative to Abe's standard of taste and false relative to Ben's? Aren't they just talking past each other, in some sense? What, exactly, does "disagreement" amount to in a relativist semantic framework? This is a key question for the relativist, because the advantage of relativism over contextualism is supposed to be precisely that the relativist secures the possibility of genuine disagreement about matters subjective.

3 What Is Disagreement, Anyway?

Before we worry about what the relativist should say about disagreement, let's think about what the *non*-relativist should say. What is it, exactly, for two parties to disagree?

3.1 Rejecting what the other accepts

The obvious thing to say is that they disagree just in case

ACCEPT/REJECT. There is a proposition that one party accepts and the other rejects.³

Perhaps it is because ACCEPT/REJECT is such an obvious answer that philosophers have not wasted much ink on the question of what it is to disagree. Surprisingly, though, this answer is wrong, or at least incomplete. This can be seen most clearly when we relativize propositional truth to parameters besides just worlds. Consider, for example, tensed propositions, which have truth values relative to world/time pairs. One such proposition

is the proposition that Joe is sitting. (Do not confuse this with the proposition that Joe is sitting now, or at any other time: the tensed proposition is, in Kaplan's terms, "temporally neutral.") If you asserted this proposition at 2 p.m. and I denied it at 3 p.m., we have not in any real sense disagreed. Your assertion concerned Joe's position at 2 p.m., while my denial concerned his position at 3 p.m.⁴ So accepting and rejecting the same proposition cannot be sufficient for genuine disagreement.

Lest anyone be tempted to save ACCEPT/REJECT by denying that propositions can be "temporally neutral," the point can be made just as well with eternal propositions (with truth values relative to worlds but not times). Just as ACCEPT/REJECT can serve as a criterion for disagreement about tensed propositions only when the acceptance and rejection take place at the same time, so it can serve as a criterion for disagreement about eternal propositions only when the acceptance and rejection take place in the same world. Consider Jane (who inhabits this world, the actual world) and June, her counterpart in another possible world. Jane asserts that Mars has two moons, and June denies this very proposition. Do they disagree? Not in any real way. Jane's assertion concerns our world, while June's concerns hers. If June lives in a world where Mars has three moons, her denial may be just as correct as Jane's assertion. (Nothing hinges here on the realist talk of worlds and counterparts. However you think of modality, it makes sense to ask whether in saying what one *would have* said, in some counterfactual situation, one would have disagreed with what one actually *did* say. That you would have rejected the proposition you actually accepted is not sufficient for an affirmative answer to this question.)

Neither of the examples we have considered so far is of the "relativist" sort. Even in standard frameworks, then, we need a richer account of disagreement than acceptance and rejection of the same proposition.

3.2 Bringing in contexts

The preceding considerations suggest that an adequate account of disagreement will have to consider not just the propositional contents of the speech acts and mental attitudes at issue, but the contexts in which they occur. So let us introduce some terminology to make this easier:

ACCURACY. An acceptance (rejection) is **accurate** just in case the proposition accepted is true (false) at the circumstance of evaluation that is relevant to the assessment of the acceptance (rejection) in its context (or at all such circumstances, if there is more than one).⁵

The circumstance of evaluation that is relevant to the assessment of an acceptance or rejection of a tensed proposition is the pair consisting of the world and time of the context of use. Thus, an assertion at 2 p.m. of the tensed proposition that Joe is sitting is accurate iff this proposition is true at $\langle @, 2 \text{ p.m.} \rangle$, while a denial at 3 p.m. of the very same proposition is accurate iff this proposition is false at $\langle @, 3 \text{ p.m.} \rangle$.⁶ Since it may be that Joe is sitting at 2 p.m. but not at 3 p.m., it is clear that an assertion and a denial of the same proposition can both be accurate.

It is common in philosophy to express this thought by saying that both “utterances” can be true. But it seems to me that utterances in the “act” sense are not properly classed as true or false at all, while utterances in the “object” sense—sentence types or tokens—have truth values only relative to contexts. (And yes, a token, as well as a type, can be used at different contexts.) So I have chosen to use a different predicate, “accurate,” for this classification of assertions and denials, considered as acts.⁷

How might we explain disagreement in terms of accuracy? A minimal strengthening of ACCEPT/REJECT would be

NOT BOTH ACCURATE. (a) There is a proposition that one party accepts and the other rejects, and (b) it is not the case that both the acceptance and the rejection are accurate.

But this is still not a sufficient condition for genuine disagreement. Suppose that at noon Mary accepts the tensed proposition that Socrates is sitting, and at midnight Tom rejects this proposition. And suppose that Socrates was sitting at both noon and midnight. Then Mary's acceptance is accurate, while Tom's rejection is inaccurate, but still there is no disagreement. For the acceptance and rejection concerned entirely different times; indeed, they *could* both have been accurate, even if in fact they were not.

This last remark suggests that what is needed is to supplement NOT BOTH ACCURATE with some modality:

CAN'T BOTH BE ACCURATE. (a) There is a proposition that one party accepts and the other rejects, and (b) the acceptance and the rejection *cannot* both be accurate.

However, we still don't have a sufficient condition for disagreement. Suppose that at noon Mary accepts the tensed proposition that the number of flies in the room is either odd or even, and at midnight Tom rejects this same tensed proposition. Then Tom's rejection *can't* be accurate (he is, after all, rejecting a necessary truth), so *a fortiori* Mary's acceptance and Tom's rejection can't *both* be accurate. Yet they do not disagree: Mary's thought concerns noon, while Tom's concerns midnight.⁸ We could avoid this problem by excluding cases where one of the two acts in question is *necessarily* inaccurate. But then we would exclude too much: in noncontingent domains, like mathematics, disagreement *always* involves an acceptance or rejection that is necessarily inaccurate.

We have been focusing so far on the *insufficiency* of the proposed criteria. But NOT BOTH ACCURATE (and hence also CAN'T BOTH BE ACCURATE) is not even a *necessary* con-

dition for disagreement. For it requires that in cases of disagreement there be a single content p that one party accepts and the other rejects. And that is too strong. If at noon Mary accepts the tensed proposition that Socrates is sitting, and at midnight Peter rejects the tensed proposition that Socrates was sitting twelve hours ago, then they disagree. Peter is not rejecting the proposition that Mary accepted, but a suitably related one. Thus NOT BOTH ACCURATE is, at best, a necessary condition for disagreement in the special case where a single proposition is at issue. To get a fully general necessary condition, we would need to say what it is for two propositions to be “suitably related” such that acceptance of one and rejection of the other constitute disagreement. This is not too hard, perhaps, for tensed propositions, but one wants a general account of disagreement that works no matter what propositional truth is relative to.

Even without looking at distinctively relativist views, then, it is difficult to give a fully general necessary and sufficient condition for disagreement. Without such a condition, we won't be able to settle definitively whether a relativist semantics can secure genuine disagreement. But I think we can still make some progress on this question by taking CAN'T BOTH BE ACCURATE as our working account of disagreement, despite its flaws. I don't have anything better to offer, and CAN'T BOTH BE ACCURATE does seem to capture the core of the idea of disagreement.

4 Relativism and accuracy

What should the relativist say about accuracy? Earlier we said that an acceptance (rejection) of p at context C is accurate iff p is true (false) at the circumstance relevant to the assessment of the acceptance (rejection) in its context. But which circumstance is that? For the temporalist, as we have seen, it is the circumstance consisting of the world and time of the context of use:

TEMPORALIST ACCURACY. An acceptance (rejection) of a proposition p at C is accurate iff p is true (false) at the circumstance $\langle w_C, t_C \rangle$, where w_C is the world of C and t_C is the time of C .

But how should we deal with the “non-standard” parameters of propositional truth posited by the relativist, like standards of taste?

4.1 Use-centric accuracy

One option is to treat these extra parameters the same way the temporalist treats the time parameter:

USE-CENTRIC ACCURACY. An acceptance (rejection) of a proposition p at a context C is accurate iff p is true (false) at the circumstance $\langle w_C, s_C \rangle$, where w_C is the world of C and s_C is the standard of taste of the speaker at C .

But this account of accuracy does not give us genuine disagreement in the cases of interest to the relativist. Suppose Abe accepts the proposition that apples are delicious, while Ben rejects it (at the same time, in the same world). According to USE-CENTRIC ACCURACY, Abe’s acceptance and Ben’s rejection can both be accurate, provided Abe and Ben have sufficiently different standards of taste. But surely NOT BOTH ACCURATE is a necessary condition for disagreement, at least when it concerns a single proposition. So on this account, Abe and Ben do not really disagree.

It is easy to miss this point, because USE-CENTRIC ACCURACY does predict that Ben is accurate in accepting the proposition that the proposition Abe accepts is false, while Abe is accurate in accepting the proposition that the proposition Ben rejects is true. (Here I assume that the proposition *that p is true* is true at a circumstance $\langle w, s \rangle$ iff p is true at $\langle w, s \rangle$, and that the proposition *that p is false* is true at $\langle w, s \rangle$ iff p is false at $\langle w, s \rangle$.) But

the fact that Ben can (accurately!) tell Abe that what he accepts is false does not show that there is any real disagreement between them, any more than the fact that I can now (accurately) assert that the tensed proposition *that it is night* is false shows that there is any real disagreement between me and someone who asserted this proposition eight hours ago.

Perhaps some relativists will be satisfied with this kind of “faux disagreement.” But I don’t think such a position accounts for *all* of the data that motivate relativist semantics (for example, the retraction data in the case of epistemic modals). Indeed, I am inclined not to call such a position a “relativist semantics” at all, since formally it is no different from what can be found in Kaplan’s “Demonstratives.” I have elsewhere suggested using the term “nonindexical contextualism” for positions like this, which take truth to depend on features of the context of use, but locate the context sensitivity in the context’s *circumstance-determinative* role rather than its *content-determinative* role (MacFarlane forthcoming b, forthcoming a).

4.2 Perspectival accuracy

In order to get a real relativism that allows genuine disagreement about matters of taste, we need “accuracy” to depend not just on the proposition at issue and the context in which it is used, but on the context in which it is *assessed* (what I have elsewhere called the “context of assessment,” by analogy with the “context of use”). That is, we need accuracy to be perspectival or “assessment-sensitive.”

PERSPECTIVAL ACCURACY. An acceptance (rejection) of a proposition p at a context C_U is accurate (as assessed from a context C_A) iff p is true (false) at the circumstance $\langle w_{C_U}, s_{C_A} \rangle$, where w_{C_U} = the world of C_U and s_{C_A} = the standard of taste of the assessor at C_A .

On this view, an acceptance of p is accurate just in case p is true at the world of the context of use and the standard of taste of the assessor at the context of assessment. The generalization to “funny” and “likely” is obvious.

To be a relativist, then, is not to relativize propositional truth to “nonstandard parameters” like standards of taste,⁹ but to adopt a certain view about how the *accuracy* of certain acts or states is to be assessed.¹⁰

What happens to disagreement when accuracy becomes perspectival? Lacking anything better, we’ll stick to CAN’T BOTH BE ACCURATE as a criterion for disagreement. But how should we apply this when accuracy becomes perspectival? The natural thing to do is to let disagreement, too, be perspectival. Thus, two parties disagree (as assessed from context C) if

CAN’T BOTH BE ACCURATE (RELATIVE TO C). (a) There is a proposition that one party accepts and the other rejects, and (b) the acceptance and the rejection *cannot* both be accurate (as assessed from C).

This account of disagreement, combined with PERSPECTIVAL ACCURACY, gives us the result we wanted in the case of Abe and Ben. ACCORDING TO PERSPECTIVAL ACCURACY, Abe’s acceptance is accurate (as assessed from C_A) only if the proposition that apples are delicious is true at $\langle w, s_{C_A} \rangle$, while Ben’s rejection is accurate (as assessed from C_A) only if this proposition is false at $\langle w, s_{C_A} \rangle$. (w here is the world in which Abe accepts and Ben rejects the proposition that apples are delicious, and s_{C_A} is the standard of taste of the assessor at C_A .) Since no proposition can be both true and false at the same circumstance of evaluation, Abe’s acceptance of the proposition that apples are delicious and Ben’s rejection of this proposition cannot both be accurate. As assessed from C_A , then, Abe and Ben disagree. It should be clear that this result will hold equally for any choice of C_A . Thus, although technically we have made disagreement perspectival, there is little cost to

this, as it will never happen that two parties disagree as assessed from one context but not as assessed from another.

5 Understanding Perspectival Accuracy

I have argued that what makes a view “relativist” is its relativization of the notion of *accuracy* to contexts of assessment. But what does it *mean* to say that a speech act or mental state is accurate “from a perspective” or “relative to a context of assessment”? The relativist needs to say something about the practical significance of claims of assessment-relative accuracy. How does it matter in practice whether a speech act or mental state is accurate relative to one context of assessment rather than another? What turns on this?

5.1 Accuracy as aim

One thing we might say about accuracy is this:

ACCURACY AS AIM. We *aim* for our speech acts and mental states to be accurate. When they are not accurate, they have, in a certain sense, “misfired.”

Making this idea out in any detail is notoriously difficult. At least in the case of speech acts, there seem to be exceptions. Some assertions—lies—aim for *inaccuracy*, although even they represent themselves as aiming at accuracy. It may be that *beliefs* necessarily aim at accuracy (cf. Williams 1973), but resolving this question requires cashing out the metaphorical talk of beliefs “aiming” at something.

Even if some suitably qualified version of ACCURACY AS AIM is correct, however, it is not clear how it could help us make sense of perspectival assessments of accuracy. For when we are talking about achieving aims, there is a privileged context of assessment that matters in a way that the others don’t—the one occupied by the person who *has* those

aims. That is the context one must keep an eye on if one is trying to understand the agent's actions in terms of her aims.

Consider two cases in which *A* asserts that *p* at C_1 . In the first case, *A*'s assertion is accurate relative to both C_1 and some different context C_2 (occupied by some other agent), while in the second case, it is accurate relative to C_1 but not C_2 . The relativist is committed to the idea that we have described two genuinely different cases here. But how, in practice, can we distinguish them? Talk of *aim* seems too agent-centered to give us much grip on the difference.¹¹

5.2 Challenge and response

I think that one can do better, at least in the case of speech acts, by talking about the *normative significances* of assertions (or denials) rather than their aims. An assertion is a particular kind of act. Normally, it involves the utterance of a sentence, but it is more than that. What more? What is one doing when one utters a sentence to make an assertion, and not, say, to make a conjecture, or practice pronunciation? One may have various communicative intentions: for example, one may intend to get one's audience to do something, or to believe something, or to believe that one believes something. But none of these intentions is a necessary condition for making an assertion. What *is* common to every case of assertion, even transparently insincere and soliloquizing assertions, is a certain kind of commitment undertaken by the asserter: a commitment to the truth of the proposition asserted.

But what is it to commit oneself to the truth of a proposition? If we are going to understand assertion as commitment to truth, we need to understand what such a commitment requires of one. What is one committed to *doing* (or refraining from doing) when one commits oneself to the truth of a proposition? What would count as honoring this

commitment, and what would count as violating it?

The best answer I have seen to this question is Brandom's (1983, 1994)—though I will modify his view in several respects, and use it to different ends. On this account, an assertion is fundamentally a move in the “game of giving and asking for reasons.” In making an assertion, one licenses others to rely on its accuracy in their actions and reasoning, and one commits oneself to vindicating its accuracy in the face of appropriate challenges. Thus, assertoric commitment is commitment to meeting all legitimate challenges to the accuracy of one's assertion, and to withdrawing the assertion (disavowing the commitment) if one cannot do so.

This account gives us new resources for understanding what disagreement looks like for the relativist. For now it is not just the asserter's assessments of accuracy that matter. When an assertion is challenged, there are always two relevant perspectives: that of the asserter and that of the challenger. We may suppose that

- one is entitled to challenge an assertion when one has good grounds for thinking that it was not accurate, and
- a successful response to such a challenge consists in a demonstration that the assertion was, in fact, accurate.¹²

This, then, is the practical significance of the classification of assertions into accurate and inaccurate:

ACCURACY AND CHALLENGES. Accuracy is the property we must show assertions to have in order to vindicate them in the face of challenges, and it is the property we must show others' assertions not to have if our challenges are to be justified.

This account can be carried over unchanged if accuracy is relativized. Making the relativitization explicit, we can see that there is work for a relativized notion of accuracy to do:

- one is entitled to challenge an assertion when one has good grounds for thinking that the assertion was not accurate (relative to the context of assessment one occupies in issuing the challenge), and
- a successful response to such a challenge consists in a demonstration that the assertion was, in fact, accurate (relative to the context of assessment one occupies in giving the response).

This account captures the distinctive phenomenology of disagreement about matters whose truth is relative. The challenger thinks (rightly) that he has absolutely compelling grounds for thinking that the assertion was not accurate. But the original asserter thinks (also rightly, from her point of view) that the challenger's grounds do nothing to call in question the accuracy of the assertion. The asserter's vindication will seem to the challenger not to show that the assertion was accurate, and the challenger will continue to press his claim. (Until the game gets boring.) Thus we have all the normative trappings of real disagreement, but without the possibility of resolution except by a relevant change in one or both parties' contexts of assessment.¹³

6 The Point of Assessment Sensitivity

This can look like a pretty silly game. Why do we play it? Assuming we do have assessment-sensitive expressions in our languages, why do we have them? What would we be missing if we did not?

Here's a thought experiment. Imagine a world in which we have no terms that work like 'delicious' or 'funny.' People never say that foods are delicious; they only say that foods taste good *to them*. People never say that jokes are funny; they only say that jokes are amusing *to them*. How would such a world be different from ours?

Clearly, there would be a lot less controversy! Compare the following:

(1a) "Apples are delicious."

"No, they aren't, they're bland!"

(1b) "Apples are very pleasing to my taste buds."

"#No, they aren't! They aren't very pleasing to my taste buds."

(2a) "That joke was funny."

"No, it wasn't. It was stupid."

(2b) "That joke was amusing to me."

"#No, it wasn't. It wasn't amusing to me."

Assessment-sensitive expressions are designed, it seems, to foster controversy, where use-sensitive expressions preclude it. But what is the point of fostering controversy in "subjective" domains, if there is no (nonrelative) truth on which both parties can converge? Why shouldn't we just talk about our own tastes, rather than ascribing subjective properties to the objects?

Perhaps the point is to bring about agreement by leading our interlocutors into relevantly different contexts of assessment. If you say "skiing is fun" and I contradict you, it is not because I think that the proposition you asserted is false as assessed by you in your current situation, with the affective attitudes you now have, but because I hope to change these attitudes. Perhaps, then, the point of using controversy-inducing assessment-sensitive vocabulary is to foster *coordination* of contexts. We have an interest in sharing

standards of taste, senses of humor, and epistemic states with those around us. The reasons are different in each case. In the case of humor, we want people to appreciate our jokes, and we want them to tell jokes we appreciate. In the case of epistemic states, it is manifestly in our interest to share a picture of the world, and to learn from others when they know things that we do not.

Controversy encourages coordination because, in general, controversy is uncomfortable. But why should controversy feel uncomfortable even when the disagreement is entirely due to differences in the interlocutors' respective contexts of assessment? One possible answer is: it just is. That's a brute psychological fact about us. Perhaps, as Allan Gibbard suggests (Gibbard 1990, 217), there is an evolutionary explanation. Assessment-sensitive expressions exploit this psychological fact about us—our tendency to treat dispute as a crisis to be resolved—to foster subjective coordination by provoking controversy. From lofty philosophical heights, the language games we play with words like 'funny' and 'likely' may seem irrational. But that is no reason to deny that we do play these games, or that they have a social purpose. If describing our use of these expressions requires relativizing the accuracy of speech acts and mental states to contexts of assessment, then this much relativism is required by our dispensation to describe the world as it is, not as we think it ought to be. That is surely a motivation that even the staunchest antirealist can embrace.

Notes

¹Here we might appeal either to the coinvestigators' *common* knowledge or to their *distributed* knowledge. *p* is common knowledge in a group just in case every member of the group knows that *p*, knows that the others know that *p*, knows that the others know that they know that *p*, and so on. *p* is distributed knowledge in a group just in case a third party who knew everything known by the individual group members

would be in a position to know that p (see Teller 1972).

²For a similar point about epistemic contextualism, see Feldman 2001, 29.

³I use the generic terms “accept” and “reject” here to cover both speech acts (assertion and denial) and mental acts or states (belief and disbelief). I will assume that one way to reject a proposition is to accept a proposition incompatible with it—where two propositions are incompatible iff there is no circumstance of evaluation at which both are true. But I will not assume that that this the only way to reject a proposition.

⁴For the useful distinction between what a claim is “about” and what it “concerns,” see Perry 1986.

⁵In most applications, context will single out one circumstance as relevant—what Kaplan calls “the circumstance of the context” (Kaplan 1989, 522), but this need not always be the case. For an application where the assumption of uniqueness fails, see MacFarlane 2003 and MacFarlane forthcoming c.

⁶Here ‘@’ denotes the actual world.

⁷An alternative would be to talk of the truth of propositions relative to *contexts*, as I do in MacFarlane 2005.

⁸Of course, they may disagree about another matter—whether there are numbers that are neither even nor odd. But they need not disagree even about this. Tom may have simply failed to apply his general arithmetical knowledge to the question at hand, or he may have no opinion about the general question.

⁹For this characterization, see, for example, Kölbel 2002; Stanley 2005; Zimmerman forthcoming; Egan forthcoming.

¹⁰Indeed, one can be a relativist in this sense without relativizing propositional truth to anything besides worlds, as I show in MacFarlane forthcoming c.

¹¹For a related discussion, see MacFarlane 2005, 331–2.

¹²Note that given our definition of ‘accurate,’ ‘is accurate’ and ‘was accurate’ are equivalent: the tense on the copula is determined grammatically by the subject, but has no semantic significance.

¹³Of course, one can—for a time—experience the same phenomenon even in “fully objective” discourses, when the two parties have very different background assumptions. What seems to one party to be an utterly compelling proof seems to the other to be devoid of force. What is the difference, in practice, between these cases and the relativist cases? One is tempted to say: in the relativist cases, the phenomenon arises from a relativity about what it is for an assertion to *be* accurate, while in the objective cases, it arises from a relativity about what counts as an adequate proof. But we were hoping to get beyond characterizing relativism in terms of accuracy, by saying something about the practical significance of accuracy classifications. The

difficulty is that what has practical significance is *warranted judgements about accuracy*, and it is hard to separate out what relativity is due to the “warranted” part and what is due to the “accurate” part.

References

- Brandom, R. (1983). Asserting. *Noûs*, 17, 637–650.
- Brandom, R. (1994). *Making It Explicit* (Cambridge: Harvard University Press).
- DeRose, K. (2004). Single scoreboard semantics. *Philosophical Studies*, 119, 1–21.
- Egan, A. (forthcoming). Epistemic modals, relativism, and assertion. *Philosophical Studies*.
- Feldman, R. (2001). Skeptical problems, contextualist solutions. *Philosophical Studies*, 103, 61–85.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge: Harvard 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 & H. Wettstein (Eds.), *Themes from Kaplan* (Oxford: Oxford University Press), (pp. 481–566).
- Kölbel, M. (2002). *Truth Without Objectivity* (London: Routledge).
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly*, 53(212), 321–336.

MacFarlane, J. (2005). Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105, 321–339.

MacFarlane, J. (forthcoming a). Nonindexical contextualism.

MacFarlane, J. (forthcoming b). Semantic minimalism and nonindexical contextualism. To appear in G. Preyer and G. Peter (eds.), *Content and Context: Essays on Semantics and Pragmatics*.

MacFarlane, J. (forthcoming c). Truth in the garden of forking paths.

Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian Society*.

Stanley, J. (2005). *Knowledge and Practical Interests* (Oxford University Press).

Teller, P. (1972). Epistemic possibility. *Philosophia*, 2, 303–320.

Williams, B. (1973). Deciding to believe. In *Problems of the Self* (Cambridge: Cambridge University Press), (pp. 136–151).

Zimmerman, A. (forthcoming). Against relativism. *Philosophical Studies*.