Epistemic Modalities and Relative Truth*

John MacFarlane†

Draft of November 13, 2003

I want to discuss a puzzle about the semantics of epistemic modals, like "It might be the case that" as it occurs in "It might be the case that Goldbach's conjecture is false."[1] I'll argue that the puzzle cannot be adequately explained on standard accounts of the semantics of epistemic modals, and that a proper solution requires relativizing utterance truth to a context of assessment, a semantic device whose utility and coherence I have defended elsewhere for future contingents (MacFarlane 2003).

1 Standard semantics for epistemic modals

Standard accounts of the semantics of epistemic modals take them to be context-sensitive expressions. The simplest such account is the following:

> **(Simple)** 'It might be the case that p' is true as uttered by S at t iff S does not know that not-p at t.[2]

On this account, 'It might be the case that p' means something like 'For all I know, p might be true,' or 'I don't know that p is not true.'

*I presented this paper to the Philosophy Department at the University of Utah on October 31, 2003, and I am grateful to the audience there for stimulating questions. I would also like to thank Kent Bach, Matt Weiner, Brian Weatherson for useful correspondence. Finally, I acknowledge the financial support of an ACLS/Andrew W. Mellon Fellowship for Junior Faculty and a Berkeley Humanities Research Fellowship.

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

[1] I will assume, as seems plausible, that the modality here *is* epistemic. I will also assume that epistemic modals contribute to the truth-conditional contents of claims rather than their illocutionary force. See Papafragou, forthcoming, for a defense of this view against the standard arguments to the contrary.

[2] This is essentially IP-1 from DeRose 1991, 584.

Proponents of the standard semantics have modified the Simple view in two directions to accommodate various counterexamples.[3] First, they have allowed a degree of flexibility as to whose epistemic situation is the relevant one for evaluating the modal claim. Thus, 'It might be the case that p' can express any of the following, depending on the speaker's intentions in making the utterance:

1. I do not know that not-p.

2. None of the parties to this conversation knows that not-p.

3. None of the authorities and experts knows that not-p.

Second, they have allowed that in some uses 'It might be the case that p' can be false even if the speaker (or relevant community) does not know that not-p, provided that the speaker (or relevant community) could *come to know* that not-p by inference from other things known, or by relevant other means of acquiring new knowledge (e.g. looking in a book). Here is what we get when we incorporate these two dimensions of flexibility into the contextualist semantics:

> **(Complex)** 'It might be the case that p' is true as uttered by S at t iff
>
> (a) no member of the relevant community knows that not-p at t, and
>
> (b) there is no relevant way by which members of the relevant community can come to know that not-p.[4]

For our purposes, the differences between (Simple) and (Complex) won't matter much. Both accounts share an essential feature: the context of utterance (including the speaker's intentions) fixes some specific epistemic standard (a particular community's knowledge plus certain relevant ways of extending this knowledge) relative to which the claim is to be evaluated. It is this feature of the standard account that will concern us in what follows.

[3] See Hacking 1967, Teller 1972, DeRose 1991.

[4] This is a slight adaptation of DeRose's account in DeRose 1991, 594, including some of his wording. DeRose's account is maximally flexible. A popular but less flexible account would restrict the community to S and the relevant ways of knowing to logical deduction. On this account, 'It might be the case that p' is true as uttered by S at t just in case p is consistent with everything S knows at t. Another plausible account would restrict the relevant ways of knowing to *obvious* inferences.

2 The puzzle

Here's the puzzle.[5] When we claim 'It might be the case that p' and later come to know that not-p (for example, by being told this by a trusted authority), we tend to *withdraw* our earlier assertion and regard what we said as false. That is, we tend to assess past assertions containing epistemic modals against our *current* epistemic situations, rather than the epistemic situations in which they were made. Consider the following dialogue:

> **Dialogue 1**
>
> *Sally:* Joe might be in Boston. (= It might be the case that Joe is in Boston.)
>
> *George:* He can't be in Boston. (= It is not the case that it might be the case that Joe is in Boston.) I saw him in the hall five minutes ago.
>
> *Sally:* Oh, then I guess I was wrong.

This is not what one would expect if the standard semantics for epistemic modals were correct. If the truth-conditional content of Sally's original claim is just that she does not (at the time) know that Joe is not in Boston, then she should not retract this claim upon learning that Joe is not in Boston. What she said is still true, even though she could not *now* truly say "Joe might be in Boston." So if the contextualist semantics is correct—the Simple version, at any rate—then we should expect the last line of the dialogue to be:

> *Sally:* Oh, okay. So he can't be in Boston. Nonetheless, when I said "Joe might be in Boston," what I said was true, and I stand by that claim.

I hope you'll agree that it would be odd and unnatural for Sally to say this.

In this respect, "might" is very different from ordinary indexicals. Compare the following parallel dialogue:

> **Dialogue 2**
>
> *On a moving train...*

[5] For a version of this puzzle, see Hawthorne, forthcoming (1: 29 n. 69). Hawthorne mentions it in passing and does not offer a solution.

Sue: We're still in Boston.

Several minutes go by…

Jim: We're not in Boston. We just passed the city limits sign.

Sue: Oh, then I guess I was wrong when I said, "We're still in Boston."

Anyone who heard this exchange would assume that Sue is horribly confused, or perhaps that she had dozed off. What Jim tells her gives her no reason at all to retract her original claim. Yet the parallel dialogue with "might" seems perfectly natural. This disanalogy between standard indexicals and "might" suggests to me that "might" is not context-sensitive in the way that the standard account says it is.

It might be thought that the Complex version of the standard account can do better here than the Simple version. According to (Complex), Sally can assert "Joe might be in Boston" with the intention that the relevant community include George, or with the intention that the relevant "ways of coming to know" include being told by George. In that case, she would be entirely right to retract her assertion, since it would have been shown to be false (on the standard semantics).

The added flexibility of (Complex) helps, but it is not enough. For one thing, in some cases the relevant community may include *just Sally*, and there may be no relevant "ways of coming to know."[6] (Complex) predicts that in these cases, at least, Sally should stand by her original claim, even after she learns that Joe is not in Boston. But it seems to me that the alternative continuation to Dialogue 1, on which Sally stands by her original claim, *always* sounds wrong. We simply don't have the practice of standing by old claims of epistemic possibility in the face of new knowledge. We *could have* had this practice, but as things are, we just don't talk this way.

Moreover, even when the relevant community does include George and there are some relevant "ways of coming to know," (Complex) seems to yield the wrong verdicts. Unless the relevant community is all-encompassing, we can imagine someone who is not in this community overhearing the conversation and assessing Sally's utterance. It seems to me that this person will take what Sally says to be false, not true:

[6]DeRose 1991 calls this the "solitary use" (596).

Dialogue 3

Jane, a stranger, is hiding in the bushes…

Sally: Joe might be in Boston.

George: Oh, really? I didn't know that.

Jane (sotto voce): Sally is wrong. I saw Joe just a few minutes ago.

According to (Complex), Jane is mistaken. What she should say (if she wants to speak truly) is:

Jane (sotto voce): Joe can't be in Boston. I saw him just a few minutes ago. Nonetheless, what Sally said is true.

But I think we'd find it very strange if Jane did say this. Maybe it is the right thing for her to say, but in that case we need an explanation of why it sounds so wrong.

Similarly, unless the relevant "ways of coming to know" are all-encompassing, we can imagine Sally learning that Joe is not in Boston in an "irrelevant" way. It seems to me that in such a case, she will still regard her previous claim as false and retract it:

Dialogue 4

Sally: Joe might be in Boston.

George: Hey, look here, in the trash can. This is George's itinerary – and he's in Iceland right now.

Sally: Oh, I guess I was wrong then. He can't be in Boston.

On the standard account—assuming that this fortuitous discovery does not count as a "relevant way of coming to know"—Sally should say instead:

Sally: Oh, then he can't be in Boston. Nonetheless, I was right when I said, "Joe might be in Boston."

But this would strike most speakers as tantamount to a contradiction.

In sum, these data about our linguistic practices in using epistemic modals do not seem to fit the standard contextualist semantics very well. Let us now consider what we might say to explain the data.

3 An error theory

If we stick with the standard semantics, we must say that Sally is *wrong* to retract her assertions in Dialogues 1 and 4, and that Jane is wrong to claim that Sally's assertion is false in Dialogue 3. To maintain this position is to claim that ordinary speakers make systematic errors about the truth conditions of sentences containing epistemic modals. So it's a kind of *error theory.* To be sure, sometimes people assert things they know to be (literally) false, for all kinds of reasons. But I can't think of any reason why Sally would retract her assertion in Dialogue 1 if she did not think it had been proven untrue. (Or rather, I can think of reasons, but they require elaborate surrounding stories, and the puzzling data do not.) So it seems to me that only an error theory offers any hope of reconciling the data with the standard semantics.

The problem with this option is that the arguments motivating the standard semantics (in its various versions) would be undercut by such an error theory. Consider one of the opening arguments of DeRose's "Epistemic Possibilities" (DeRose 1991). Jane's husband has gone in for some tests, which she knows might reveal conclusively that he has cancer. Jane knows that the tests have already been performed and that the doctor can rule out cancer if they are negative. In fact, the tests are negative, but Jane hasn't yet talked with the doctor and so doesn't know the results. In this situation, DeRose notes, Jane might say "It is possible that my husband has cancer," while the doctor might say "It is not possible that he has cancer." DeRose claims that both are speaking truly, and it is this that forces a contextualist semantics for epistemic possibility (582–3).

But if we allow the possibility that ordinary speakers (like Jane and the doctor) are mistaken about the truth-conditions of sentences containing epistemic modals, then we can't conclude anything about the actual truth-conditions of these sentences from data about how they use them, and DeRose's argument for contextualism is blocked. DeRose is, of course, fully aware of this. He simply presupposes that speakers are not semantically ignorant in the way the error theory posits:

> I am assuming . . . that if the analysis were correct, ordinary speakers like Jane would judge that things were possible, not possible, very likely possible, very likely not possible, etc. when the conditions of the analysis seemed to them to be satisfied, not satisfied, very likely satisfied, very

> likely not satisfied, etc. I am assuming, not that ordinary speakers have explicit knowledge of the truth conditions of "It is possible that P," but rather that they know when and how to use the expression, and that one job of analytic philosophers is to make explicit the conditions of such expressions that ordinary speakers show implicit knowledge of in their use of the expressions. (585 n. 4)

It should be clear, then, that the kind of error theory we have been contemplating as a way to *save* the standard account would in fact undercut the very arguments that are supposed to *motivate* that account.

4 A universalist semantics

If we reject the standard semantics, we need an alternative—one that explains the data in the puzzle. One possibility is to replace (Complex)'s contextualist talk of "the relevant community" and "relevant ways of coming to know" with straightforward universal quantifiers over all knowers and all possible ways of coming to know. This yields a "universalist" semantics:

> **(Universal)** 'It might be the case that p' is true as uttered by S at t iff
>
> (a) no one knows that not-p at t, and
>
> (b) there is no way in which anyone can come to know that not-p.

On this semantics, Sally is right to retract her assertions in Dialogues 1 and 4, and Jane is right to claim that Sally's assertion is false in Dialogue 3.

Note that the move to (Universal) is only going to explain Sally's retractions if we suppose that she *knows* the truth conditions of "might," as given by (Universal). If she were in error about the semantics of "might," thinking, for example, that they were as given by (Simple), then why would she concede that she had been wrong? She certainly hasn't learned anything new about the semantics of "might" in the course of the dialogue. So we'll assume that Sally has at least an implicit knowledge of (Universal).

But if "Joe might be in Boston" means what (Universal) says it does, and Sally knows this, why would she have said it in the first place? This needs explaining just as much as her subsequent

retraction. And it is hard to see how we *could* explain it if (Universal) is correct. According to (Universal), the truth-conditional content of Sally's claim is that (a) no one knows that Joe is not in Boston, and (b) there is no way in which anyone can come to know that Joe is not in Boston. This is an incredibly strong claim, and Sally certainly does not have grounds for thinking that it is true. In fact, if she is a reasonable person she will conclude that it is almost certainly false. Why, then, would she make this claim? We can assume that she is not trying to mislead her audience: after all, they are as well placed as she is to see that her claim, construed as (Universal) construes it, is false.

The possibility remains that Sally is saying something patently false in order to convey something true, as one might say "He can throw a ball a million miles" in order to convey the thought that he can throw a ball a really long way. Perhaps Sally is trying to convey the thought that *she* doesn't know that Joe isn't in Boston. But it is hard to see why she would say something equivalent to "No one knows that Joe is in Boston, and there's no way anyone can come to know this" in order to convey "*I* don't know that Joe is in Boston."

The problem with (Universal) is that it goes too far in divorcing the truth conditions of "It might be the case that p" from the epistemic situation of the speaker. It makes "It might be the case that p" tantamount to "There is no way to come to know that not-p," or "p is not falsifiable." That means that "It might be the case that p and it might be the case that not-p" will only be true when p is in principle undecidable. There is no doubt a use for an operator that works this way, but this is simply not the way our ordinary epistemic modal talk works.

Suppose Anne and Jim are trying to determine whether a complex formula of propositional calculus is a tautology. The formula has so many variables that it would take years to grind through a truth table.

Dialogue 5

Anne: This formula might be a tautology, and it might not be.

Jim: But it's decidable whether a formula of propositional calculus is a tautology.

Anne: Right. In principle, we could find out the answer, and I hope we will find it soon. But all we know now is that it might be a tautology, and it might not be.

Anne does not take Jim's observation to be a reason for thinking her assertion of epistemic possibility to be false. But according to (Universal), she ought to. Perhaps (Simple) was too subject-centered in its assignment of truth conditions to sentences containing epistemic modals, but (Universal) goes too far in the other direction. It is the desire for a middle ground that motivates a semantics like (Complex), which allows the degree of subject-centeredness to vary with the context of use, but as we have seen, (Complex) does not adequately handle the data scouted in section 2. A new approach is needed.

5 An assessment-sensitive semantics

The solution I have in mind requires a departure from standard semantic ideas. In standard semantics, truth for sentences is relativized to contexts of use, in order to accommodate indexicals, tense, and other context-sensitive elements. But truth for *utterances* is absolute.[7] A particular assertion may be true or false or (perhaps) neither, but whatever truth value it has, it has absolutely, independent of the perspective from which it is assessed. I think that it is time to revisit this standard assumption. The semantics of certain expressions cannot be adequately represented, I would urge, unless we allow utterance truth to vary with the context in which the utterance is assessed: what I call the *context of assessment*, by analogy with the context of use. Accordingly, sentence truth must be relativized not just to a context of use, but to a context of assessment.

It is important to see the difference between relativizing truth to a context of assessment and relativizing truth to a circumstance of evaluation. The latter is an essential part of standard semantics. We could not do semantics for temporal or modal operators without relativizing truth to a time and world. But this kind of relativity is fully compatible with the absoluteness of utterance truth. Suppose I assert "Bush is president." This sentence (or, if you like, the proposition it expresses

[7] "Utterance" is ambiguous between the act of utter*ing* and the object thereby produced; I use it in the "act" sense throughout. Thus a "sentence token"—a written inscription or acoustic wave pattern in the air—is not an utterance in my sense, and the fact that a sign on an office door saying "Back in an hour" can be accurate at one time and inaccurate at another is not a counterexample to the standard assumption that utterance truth is absolute (*pace* Percival 1994, 205). Nor is the fact that an answering message "I am not here now" may give correct information to one caller and incorrect information to another (cf. Predelli 1998). In that case the owner of the machine has managed to make several distinct utterances, over a period of time, with one vocalization.

in context) is true relative to certain circumstances of evaluation, false relative to others. Yet my *utterance* has its truth value absolutely. It is true *simpliciter*, because its sentential vehicle (or, if you like, its propositional content) is true relative to the circumstances determined by its context of use. Even someone assessing my assertion from a possible world or time at which Bush is not president should agree that it is true, despite the fact that an assertion of the same sentence (with the same propositional content) would not express a truth in her mouth. Thus, relativizing truth to circumstances of evaluation does not entail relativizing truth to contexts of assessment. Only the latter requires us to give up the idea that assertions have their truth values absolutely.[8]

In this modified framework, sentences can be context-sensitive in two different ways. If their truth values vary with the context of use (keeping the context of assessment fixed), they are *use-sensitive*. If their truth values vary with the context of assessment (keeping the context of use fixed), they are *assessment-sensitive*. (Of course, a given sentence may be both use-sensitive and assessment-sensitive.) I propose that sentences containing epistemic modals are *assessment-sensitive*. Sally's first claim in Dialogue 1 in true relative to one context of assessment (the one she occupies when she makes the claim), false relative to another (the one she occupies after being told that Joe was not in Boston). Similarly, Sally's claim in Dialogue 3 is true relative to the context of assessment she occupies, but false relative to the context of assessment Jane occupies. My proposal, then, is that something like the following semantics is correct:[9]

> **(Relative)** 'It might be the case that p' is true as uttered by S at t_0 and assessed by A at t_1 iff at t_1 A does not know that not-p.

I envision two sorts of objection to this proposal. First, that it is ad hoc. Good scientific practice dictates that we make central modifications to our theories only when they have great and wide-ranging explanatory value. Surely it is not a good idea to make structural changes in our semantic framework just to accommodate epistemic modals.

[8]For further discussion of this distinction, see my manuscript "Three Grades of Truth Relativity" and MacFarlane 2003, 327–8.

[9]For expository simplicity, I have just given truth-conditions for a standalone sentence; a more adequate treatment would need to clarify how "might" behaves when embedded under tense and modal operators, and under modifiers like "For all I know…." For a more detailed framework for relative-truth semantics, see "Three Grades of Truth Relativity."

To this objection I have two replies. First, I believe that assessment sensitivity is not limited to epistemic modals. I have argued elsewhere that assessment sensitivity is the key to an adequate solution to the problem of future contingents, an adequate semantics for knowledge-attributing sentences, and an adequate account of what David Lewis called "accommodation."[10] I expect that there will be other applications as well. So the modification I propose is not tailor-made for a single use, but has much wider application.

Second, the structural changes that are required are less radical than one might think. No one should balk at the very *idea* of a context of assessment. Just as the context of use is the concrete situation in which a speech act is produced, so a context of assessment is a concrete situation in which a speech act is *assessed*. Surely there is nothing unintelligible about that. Nor is any harm done *just* by relativizing sentence truth to a context of assessment as well as a context of use. The controversial question is one that arises *within* this framework: namely, are there any sentences whose truth values *vary* with the context of assessment? If the answer is no, we may carry on within the modified framework just as before: the extra parameter for "context of assessment" will be an idle wheel, but a harmless one. Even if the answer is yes, existing accounts of the semantics of sentences that are *not* assessment-sensitive can be carried over virtually unchanged. The relativism I advocate is local, not global.

The second objection I expect is that relativized utterance truth is *incoherent*. What on earth could it *mean* to say that my utterance is true as assessed by me, now, but false as assessed by me, later; or true as assessed by me, but false as assessed by you? One often hears confused undergraduates using "p is true for Sarah" to mean, in effect, "Sarah takes p to be true." But that is a blatant conflation of truth with taking-true. How, then, should we understand "true at a context of assessment," if not in this confused way? What do I mean when I say that an utterance is true relative to one context of assessment, false relative to another?

I do not propose to answer this question by *defining* "truth at a context of assessment." I think the project of defining truth (not truth for a particular language, but *truth*) is pretty hopeless.[11] Instead, I

[10] See MacFarlane 2003 and my unpublished manuscripts "Three Grades of Truth Relativity" and "A Relativist Semantics for 'S knows that p'."

[11] See Davidson 1997.

want to explicate the notion of "truth at a context of assessment" by exhibiting the role it can play in a broader theory of meaning: a systematic account of the literal significances of speech acts that can be made using arbitrary sentences of a language. Michael Dummett (1959) took a similar approach when he argued that the notion of truth must be understood in relation to the practice of assertion: roughly, truth is what we aim at in making assertions. Obviously I can't say *this*, because on my account assertions don't in general *have* absolute truth values. What could it *be* to aim at truth, if my utterance can be true relative to some contexts of assessment and false relative to others? That would be like aiming to walk *farther away*—not farther away from here or there or this or that, but just *farther away* simpliciter. For essentially this reason, Gareth Evans (1985) rejected the idea that utterance truth could be relative.[12]

Perhaps instead we should reject the idea that truth is the normative aim of assertion. I have always found this idea obscure and of dubious utility. An account of assertion ought to tell us what one is *doing* in asserting something—not what one ought to be aiming at, or when one's speech act is "correct" *qua* speech act of that type. After all, if we wanted to know what it was to *bunt* in baseball, we would hardly be satisfied with a story about when we ought to bunt, or what we should aim at in bunting.

So, what *is* one doing when one asserts a sentence (as opposed to uttering it to practice pronunciation, or using it to ask a question)? Here is a tried and true answer:

> **(Assertion-1)** To assert a sentence S is (*inter alia*) to commit oneself to the truth of S (relative to the context of use).[13]

(Assertion-1) tells us what is accomplished by asserting a sentence, even if the assertion is insincere, deceptive, reckless, or otherwise defective. But what is it to commit oneself to the truth of a sentence (relative to a context of use)? Some philosophers seem happy to take "commitment to truth" as basic. But I agree with Belnap and Perloff that deontic constructions must take agentive complements, at

[12]Note that Evans talks of "correctness" rather than "truth" for utterances, but the difference is merely terminological. See the extended discussion of Evans' objection in section VI of MacFarlane 2003.

[13]For the idea of assertion as commitment to truth, see e.g. Searle 1979, 12. The "*inter alia*" leaves room for the possibility that asserting a sentence involves more than simply committing oneself to its truth. Plausibly, the commitment must be undertaken publicly, by means of an overt utterance; perhaps there are other conditions as well.

the level of logical form (Belnap et al. 2001, 13). We can make sense of "being committed to Al Gore," but only as meaning something like "being committed to *working for* (or perhaps *supporting*) Al Gore." Similarly, when we say that Susan is committed to a house being built on this spot, we might mean that she is committed to building a house on this spot, or to seeing to it that others build one. When no obvious agentive complement presents itself, we can't make any sense of deontic constructions at all. What would it mean, for example, to be committed to the color of the sky?

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

These observations suggest an answer to our question. What I have committed myself to doing, in asserting that Jake is in Boston, is seeing to it that I vindicate my claim if it is challenged.[14] There may be no specific sanction for failing to follow through on this commitment. But if I fail too blatantly or too frequently, others may stop treating me as a being that is capable of undertaking this kind of commitment. They may still take my utterances as expressions of my beliefs, as we take a dog's excited tail wagging as an expression of its psychological state.[15] They may even regard my utterances, if found to be reliable, as useful bits of information. But they will be treating me as a measuring instrument, not as an asserter. They will not take me to be *committing myself* to the truth of my utterances.

If this is the right way to think of "commitment to the truth of a sentence," then we may replace

[14]For this way of looking at the commitment undertaken in asserting as a "conditional task responsibility" to vindicate a claim when it is challenged, see Brandom 1983 and chapter 3 of Brandom 1994. I do not develop the idea in quite the same way as Brandom, but I am much indebted to his work.

[15]Note that this is a less exacting notion of "expressing a belief" than the one used by many speech-act theorists: see Bach and Harnish 1979, 15.

(Assertion-1) with the following:

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

A few points of clarification are in order.

1. Not every challenge counts as "appropriate." A mathematics professor who has asserted that it is impossible to trisect an angle using a compass and straightedge need not respond to every crackpot counterexample she receives in her e-mail. Whether a challenge counts as "appropriate" will depend on what was asserted, who asserted it, who is challenging it, what grounds for challenge have been offered, and the dialectical context. I will not attempt to work out criteria for appropriateness here, but it seems to me that we do distinguish appropriate from inappropriate challenges in our practice.

2. In speaking of "grounds for the truth of S (relative to U)," I do not want to imply that the grounds must be explicitly metalinguistic and semantic. "It had gills" can count as grounds for the truth of the sentence "It was a fish." In general, an utterance of sentence S_1 at context C_1 counts as grounds for the truth of a sentence S_2 relative to C_2 iff (a) S_1 is true at C_1 and (b) the truth of S_1 at C_1 entails, or is evidence for, the truth of S_2 at C_2.

3. What grounds count as "adequate" will depend on a number of factors, including the nature of the challenge and what the challenger knows. In response to some challenges, weak inductive evidence might be adequate; for others, demonstrative proof might be required. In addition, grounds are adequate only if the challenger is in a position to recognize that they *are* grounds, or would be in a position to recognize this if she had an explicit understanding of the semantics of the relevant sentences. Thus, an utterance of "The left and right footprints are of different sizes" can count as grounds for the truth of "The butler did it" if the challenger knows that the butler's feet are of different sizes, but not otherwise.

4. I leave it open when (if ever) it is appropriate to defer the responsibility for responding to a challenge to another asserter. This question is connected with delicate issues about the epistemology of testimony. Certainly it is not appropriate to defer to a speaker whom one has good reason to think is untrustworthy, or whose claim one has good reason to think is false. But when these defeating conditions are absent, it may be appropriate to defer instead of responding directly to a challenge.[16]

5. Withdrawal may be formal ("I take that back") or informal ("So I was wrong"). In some cases it may not need to be made explicit at all, as there is an expectation that an assertion proven to be false will be withdrawn. In those cases, one must speak up if one wishes to remain committed to the claim, or one will be presumed to have withdrawn it.

Note that as far as (Assertion-2) goes, a liar is breaking none of the norms of assertion, provided she abides by the conditional task-responsibility to vindicate her claim when it is appropriately challenged, and withdraws her assertion when it is shown to be false. Here the present account diverges from the "knowledge account of assertion," which takes both lies and unintentional falsehoods to violate the constitutive norm of assertion, "assert only what you know."[17] I think this is a point in its favor. Lying is not like cheating in baseball; it is more like stealing a base, which is a legal move in the game, but one that can be hard to get away with. Telling an unintentional falsehood is also not like cheating; it is more like an error in baseball.[18]

For present purposes, however, the salient advantage of (Assertion-2) is that, unlike accounts that talk of "aiming at truth," it can easily be modified to accommodate truth doubly relativized to a context of use and a context of assessment:

> **(Assertion-3)** To assert a sentence S (at a context U) is (*inter alia*) to commit oneself to providing adequate grounds for the truth of S (relative to U and one's current context of assessment)

[16] See Brandom 1994, 174–7.

[17] See Williamson 1996 and ch. 11 of Williamson 2000.

[18] A full defense of the account of assertion offered here would need to show how it can handle the data that seem to tell in favor of the knowledge account. It seems to me that (Assertion-2) can explain the infelicity of asserting "p, but I don't know that p" just as well as the knowledge account of assertion. In asserting this sentence, one is committing oneself to vindicating p if it is challenged, while conceding that one is not in a position to meet all appropriate challenges. But a fuller discussion of this issue must wait for another occasion.

> in response to any appropriate challenge, or (when appropriate) to defer this responsibility to another asserter on whose testimony one is relying. One can escape this commitment only by withdrawing the assertion.

On this account, when an assertion is challenged, the asserter must respond by providing adequate grounds for the truth of her utterance relative to her *current* context of assessment—the context she is in at the time of the challenge—even if this is very different from the context in which the assertion was made.[19]

(Assertion-3) allows us to see what it would be like to assert an assessment-sensitive sentence. For purposes of illustration, consider a word "noy" that works like "now," except that it refers to the time of assessment instead of the time of utterance. If Abe asserts "I am sitting now" and Ben asserts "I am sitting noy," what is the difference between their respective assertoric commitments? The difference is that when Abe's assertion is challenged, Abe is committed to showing that he was sitting at the time he made the assertion (that is, at the context of use), while when Ben's assertion is challenged, Ben is committed to showing that he is sitting at the time of the *challenge* (that is, at his current context of assessment). So Ben can be forced to withdraw his assertion if challenged at a later time when he is standing, while Abe's assertion can stand as long as he can show that he was sitting when he made it. Admittedly, the commitment Ben undertakes in asserting "I am sitting noy" is an odd one—one that it is hard to imagine a reason for undertaking. It is not surprising that there are no words in natural languages that work like "noy." But although Ben's assertoric commitment is odd, it is at least intelligible. There is nothing *incoherent* about committing oneself, for all future times t, to showing that one is standing at t if one's earlier assertion of "I am sitting noy" is challenged at t.[20]

In response to worries about the coherence of relativizing truth to a context of assessment, I have shown how assessment-relative truth interacts with the practical norms governing assertion. This is

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

[20] For a less contrived example, see the discussion of future contingents in MacFarlane 2003 and section 6 of "Three Grades of Truth Relativity."

enough to give a *significance* to semantic proposals, like (Relative), that use this relativized notion of truth. The difference between assessment-sensitive and assessment-insensitive semantics will come out as a difference in norms for responding to challenges and retracting assertions. Thus we can decide between (Relative), (Simple), (Complex), and other semantic proposals by considering when a speaker is obliged to withdraw an assertion in response to a challenge.

6 Comparing the three accounts

With this in mind, let us revisit Dialogue 1:

> **Dialogue 1**
>
> *Sally:* Joe might be in Boston. (t_1)
>
> *George:* He can't be in Boston. I saw him in the hall five minutes ago.
>
> *Sally:* Oh, then I guess I was wrong. (t_2)

For present purposes, we can idealize contexts of assessment as assessor/time pairs, e.g. "Sally/t_1" or "George/t_2." Let us also assume that George belongs to the community of knowers that is relevant at the context of Sally's original assertion (at t_1). Then Table 1 displays the truth values of Sally's first assertion, relative to various contexts of assessment. Note that on (Simple) and (Complex), Sally's utterance has the same truth value across all four contexts of assessment, while on (Relative), its truth value varies with the context of assessment.

Table 1: The three accounts compared.

Proposal	Sally/t_1	George/t_1	Sally/t_2	George/t_2
(Simple)	**True**	**True**	**True**	**True**
(Complex)	**False**	**False**	**False**	**False**
(Relative)	**True**	**False**	**False**	**False**

By appealing to (Assertion-3), we can extract the practical upshots of these competing semantic proposals:

- According to (Simple), Sally ought not withdraw her assertion in the face of George's challenge. The new information George offers is irrelevant to the truth of her original claim, which depends only on her epistemic state at the time of utterance. She may come to believe that it cannot be the case that George was in Boston at t_1, but this new belief is consistent with the truth of her original assertion.

- According to (Complex), Sally is *right* to withdraw her assertion at t_2. The fact that George knew (at t_1) that Joe was not in Boston is enough to make her assertion false, even though *she* didn't know that Joe was not in Boston when she made the assertion.

- (Relative) agrees with (Complex) that Sally should withdraw her assertion in Dialogue 1. But it gives a different reason. According to (Relative), Sally should withdraw her assertion not because she comes to see that *George* knew that Joe was not in Boston, but because *she* comes to know that Joe was not in Boston.

Which of these accounts is right? Which makes the best sense of our practice in using epistemic modals? I think we have plenty of reason to reject (Simple). If (Simple) is correct, then instead of withdrawing her assertion, Sally ought to stand by it, perhaps reiterating it in a way that makes explicit its context-dependence:

> Sally: Oh, so he can't be in Boston. Still, what I said before was true; I was only committing myself to the claim that my knowledge at that time did not exclude Joe's being in Boston—that *for all I knew then*, Joe might have been in Boston.

However, this isn't what a normal speaker would do. So, unless normal speakers are *systematically* too lazy to reformulate their epistemic modal claims in this way, or too ignorant to see that they *can* do so, (Simple) doesn't seem to capture our actual practice in using epistemic modals. Indeed, cases like this were precisely what motivated the move from (Simple) to (Complex).

Both (Complex) and (Relative) explain why normal speakers tend to withdraw their epistemic modal assertions, and take them to have been false, when they acquire relevant new knowledge. But

they give rather different explanations.[21] According to (Relative), they withdraw their claims because their contexts of assessments have changed, and the claims are assessment-sensitive. According to (Complex), on the other hand, they withdraw their claims because they discover something they didn't know about the context of *use*: for example, that someone in the relevant community knew something they didn't, or that there was a relevant "way of knowing" by which this piece of knowledge could be acquired. In the case of Dialogue 1, (Complex) says that Sally should withdraw her assertion because she sees that at the time she made it, *George* knew that Joe was not in Boston, whereas (Relative) says that she should withdraw her assertion because *she* knows *now* that Joe was not in Boston. Is there data that can help us choose between these alternative explanations?

I think there is. If (Complex) is right, one can challenge Sally's assertion by merely raising the possibility that *George* knew Joe's whereabouts. If Sally can't eliminate this possibility, she can't show that her claim is true, and she must withdraw it:

Dialogue 6

Sally: Joe might be in Boston.

Jim: How do you know that George doesn't know that Joe isn't in Boston? He might have seen Joe recently.

Sally: Good point. I have no idea whether George knows. I retract my claim.

To my ear, this dialogue sounds very strange. Whether George knows Joe's whereabouts seems irrelevant to the truth of Sally's claim. This suggests to me that (Complex) gives the wrong *explanation* of why Sally is right to withdraw her assertion in Dialogue 1. She should withdraw her assertion not because she comes to know that *George knows* that Joe isn't in Boston, but because she comes to know *herself* that Joe isn't in Boston. But how can her new knowledge (at t_2) matter for the truth of her earlier claim (at t_1)? It can matter if (Relative) is correct.

If this is right, then even cases like Dialogue 1—precisely the kind of cases used to motivate (Complex) over (Simple)—are better explained by (Relative). Indeed, I believe that almost *all* of

[21] DeRose would regard this as important: he says that the cases he considers support his semantic hypothesis "not because the hypothesis yields the right results in the cases, but because the hypothesis provides an intuitively correct *explanation* for why certain statements involving epistemic possibility are true, false, appropriately said, not appropriately said, etc." (DeRose 1991: 594, emphasis added).

the counterexamples against (Simple) that are used to motivate (Complex) can be explained equally well (or better) by (Relative). Consider how these counterexamples work. Generally we are asked to assess an assertion of "It might be the case that p." We are told that although the speaker did not know whether p, her companions knew that not-p, or that although the speaker did not figure it out, she could easily have deduced that not-p. After being given this information, we are asked to agree that the speaker's assertion "It might be the case that p" was false, despite the fact that she did not know that not-p. These intuitions are supposed to motivate the epicycles that distinguish (Complex) from (Simple). Notice, however, that these stories invariably give *us* (the readers) knowledge that not-p (in-the-story knowledge, anyway!) and thus put us in a context of assessment relative to which the speaker's assertion is false, according to (Relative). So our reactions to the stories are equally well explained by the hypothesis that (Relative) is the correct semantics for epistemic modals. And, as I have argued, once we move away from Dialogue 1 and consider cases (like Dialogue 3 and Dialogue 4) in which (Complex) and (Relative) disagree in their assignments of truth values, it is (Relative) that seems to give the correct verdicts.[22]

Thus, the case for (Complex) is only compelling when (Relative) is not on the table. The problem is that (Relative) usually doesn't make it to the table, because standard semantic frameworks do not make room for it. I have suggested that this is a mistake: relative truth semantics is neither ad hoc nor incoherent. Given (Assertion-3), it is a broadly empirical question whether our practice in using epistemic modals is more closely captured by (Simple), (Complex) or (Relative). The data, taken at face value, seems to me to support (Relative) over the other two accounts. Of course, the data need not be taken at face value: it may reflect systematic error on the part of speakers, or it may have an unobvious pragmatic explanation. In the absence of a credible alternative explanation of the data along these lines, however, I think we have good reason to take (Relative) seriously—and learn to live with relative truth.

[22]There remain the cases where it seems right for Jane to say "I don't know whether it's possible that John has cancer"—where she *does* know that she doesn't know whether he has cancer (DeRose 1991: 584–5). Dealing with these on (Relative) would require explaining how "know" works with an assessment-sensitive complement.

7 *Why* epistemic modals are assessment-sensitive

I've argued that epistemic modals are assessment-sensitive. Is that a brute contingent fact about English (and presumably other natural languages), or something that has a deeper explanation?

First, let's consider why we use epistemic modals: what we use them *for*. It seems to me that we use them to indicate areas of ignorance. We have a variety of interests in doing this. Sometimes we want to point out ignorance as a prophylactic against rash action: for example, if we see someone acting on the assumption that Joe is in Los Angeles, we might warn "Joe might be in Boston." But sometimes we want to point out ignorance in order to spur further inquiry. Suppose we want to know how Red Sox fans are reacting to their teams' loss in the playoffs. How are we going to find out? "Joe might be in Boston," you say, as a way of suggesting that we try calling his cell phone.

Because epistemic modals have this function of indicating ignorance, an invariant semantics like (Universal) would clearly be inappropriate. Because our epistemic states are always changing, ignorance is context-relative. So the semantics of epistemic modals ought to be context-sensitive, too. But in a relative-truth framework, there are two ways in which a sentence can be context-sensitive: it can be use-sensitive or assessment-sensitive. Is there any way to explain why epistemic modals should be assessment-sensitive rather than use-sensitive?

It is useful here to consider the difference between indicating areas of ignorance using "It might be the case that p" and doing so using "I don't know that not-p" (or "We don't know, and can't come to know in ordinary ways, that not-p"). Suppose that in Dialogue 1, Sally had said "I don't know that Joe is not in Boston" instead of "Joe might be in Boston." The dialogue would have continued as follows:

> **Dialogue 7**
>
> *Sally:* I don't know that Joe is not in Boston.
>
> *George:* He can't be in Boston. I saw him in the hall five minutes ago.
>
> *Sally:* Oh, thanks for telling me. I didn't know before, but now I know that Joe is not in Boston.

Sally would not retract her original claim or take it to have been shown false. Instead, she would reiterate it, reformulating it in the past tense, so that it is clear that it does not contradict her claim

"now I know that Joe is not in Boston."

So there seems to be an important difference between two ways we have of indicating areas of ignorance: "It might be the case that p" and "I/we don't know that not-p." The latter is clearly use-sensitive, while the former is not. Why would it be useful for us to have these two different ignorance-indicating devices in our linguistic arsenal?

Here is a thought. Negated knowledge claims—"I/we don't know that not-p"—indicate an area of ignorance, but they do so indirectly, by telling us something about the epistemic status of a person (or group of persons) at a time. Suppose we want to store this information. We need to keep track of *who* it is that is claimed to be ignorant, and *when*. By contrast, the *only* function of epistemic modal claims is to indicate a present area of ignorance. If someone says "It might be the case that p," and we want to store this information, we don't need to keep track of a person or a time. We just need to add p to our list of propositions that our knowledge leaves open.

When at some later time we come to know that not-p, we must subtract p from this list. Since we haven't been storing information about the persons and times who are ignorant, we must reckon all assertions of "It might be the case that p," whenever and by whomever they were made, as false. And we must require all such assertions to be retracted, so they won't mislead anyone. The result is that these claims are assessment-sensitive. But perhaps that isn't a big price to pay for the increased efficiency of being able to represent areas of ignorance directly, without representing detailed information about the epistemic states of people at times.

We want to have it both ways. On the one hand, we want a way of representing areas of ignorance, and since ignorance is person- and time-relative, we need a locution that is context-sensitive. On the other hand, we want to represent areas of ignorance in an efficient way that abstracts from the messy details of "who knew and when." To get all this, we need a locution that is context-sensitive but not use-sensitive. We need a locution that is assessment-sensitive.

References

Bach, K. and R. M. Harnish (1979). *Linguistic Communication and Speech Acts*. Cambridge: MIT Press.

Belnap, N., M. Perloff, and M. Xu (2001). *Facing the Future: Agents and Choices in Our Indeterministic World*. Oxford: Oxford University Press.

Brandom, R. (1983). Asserting. *Noûs 17*, 637–650.

Brandom, R. (1994). *Making It Explicit*. Cambridge: Harvard University Press.

Davidson, D. (1997). The folly of trying to define truth. *Journal of Philosophy 94*, 263–278.

DeRose, K. (1991). Epistemic possibilities. *Philosophical Review 100*(4), 581–605.

DeRose, K. (1998). Simple 'might's, indicative possibilities and the open future. *Philosophical Quarterly 48*(190), 67–82.

Dummett, M. (1959). Truth. *Proceedings of the Aristotelian Society n.s. 59*, 141–62.

Evans, G. (1985). Does tense logic rest upon a mistake? In *Collected Papers*, pp. 343–363. Oxford: Oxford University Press.

Hacking, I. (1967). Possibility. *Philosophical Review 76*, 143–168.

Hawthorne, J. (forthcoming). *Knowledge and Lotteries*. Oxford: Oxford University Press.

MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly 53*(212), 321–336.

Papafragou, A. (forthcoming). Epistemic modality and truth conditions. In A. Klinge and H. Müller (Eds.), *Perspectives on Modality (working title)*. Amsterdam: Benjamins.

Percival, P. (1994). Absolute truth. *Proceedings of the Aristotelian Society 94*, 189–213.

Predelli, S. (1998). I am not here now. *Analysis 58*, 107–115.

Searle, J. R. (1979). *Expression and Meaning*. Cambridge: Cambridge University Press.

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

Williamson, T. (1996). Knowing and asserting. *Philosophical Review 105*, 489–523.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.