Epistemic Modals Are Assessment-Sensitive*

John MacFarlane

Version of May 9, 2008

1 Introduction

By "epistemic modals," I mean epistemic uses of modal words: adverbs like "necessarily," "possibly," and "probably," adjectives like "necessary," "possible," and "probable," and auxiliaries like "might," "may," "must," and "could." It is hard to say exactly what makes a word *modal*, or what makes a use of a modal *epistemic*, without begging the questions that will be our concern below, but some examples should get the idea across. If I say "Goldbach's conjecture might be true, and it might be false," I am not endorsing the Cartesian view that God could have made the truths of arithmetic come out differently. I make the claim not because I believe in the metaphysical contingency of mathematics, but because I know that Goldbach's conjecture has not yet been proved or refuted. Similarly, if I say "Joe can't be running," I am not saying that Joe's constitution prohibits him from running, or that Joe is essentially a non-runner, or that Joe isn't allowed to run. My basis for making the claim may be nothing more than that I see Joe's running shoes hanging on a hook.

*I presented earlier versions of this paper at the University of Utah (2003), the University of Chicago (2005), Ohio State (2005), the Arché Center at St. Andrews (2005), UC Santa Cruz (2005), the Eastern Division APA meeting in New York (2005), Berkeley's Logic Colloquium (2006), and the University of Michigan Linguistics and Philosophy Workshop (2006). I am grateful to all these audiences for stimulating questions. I would particularly like to thank Kent Bach, Chris Barker, Fabrizio Cariani, Richard Dietz, Branden Fitelson, David Hunter, Graham Priest, Brian Weatherson, Matt Weiner, Seth Yalcin, and two anonymous referees for useful comments. Finally, I acknowledge the financial support of an ACLS/Andrew W. Mellon Fellowship for Junior Faculty and a Berkeley Humanities Research Fellowship. Readers of Egan et al. 2005 will notice substantial overlap in our conclusions and arguments. We arrived at them independently (though I was inspired by Hawthorne 2004, 27 n. 68, which I read in draft in summer 2003, and they by MacFarlane 2003). I shared with them an ancestor of this paper while they were writing theirs.

Clearly, epistemic modals have something to do with knowledge. But knowledge presupposes a knower or knowers. So, one ought to ask, *whose* knowledge is relevant to the truth of claims made using epistemic modals?

It is tempting to answer: *the speaker's*. On the resulting view, which I will call SOLIPSISTIC CONTEXTUALISM, "Joe might be running" expresses a truth just in case what the speaker knows does not rule out that Joe is running, and "Joe must be running" expresses a truth just in case what the speaker knows rules out that Joe is not running. For present purposes, we can leave the notion of "ruling out" schematic: we need not decide, for instance, whether knowledge that p rules out everything logically inconsistent with p. Our discussion of Solipsistic Contextualism and its variants will turn only on *whose* knowledge is at stake, not on what "ruling out" consists in. Hence we will regard theories that understand epistemic modals as quantifiers over "epistemically possible worlds" as versions of Solipsistic Contextualism, provided they take the relevant set of worlds (together with an ordering, perhaps) as determined by the *speaker's* knowledge or evidence.[1]

Solipsistic Contextualism promises to explain two facts about epistemic modals that would otherwise seem quite puzzling. First, it explains why we are normally prepared to make epistemic possibility claims on the basis of our own ignorance. If someone asks me whether Joe is in Boston, it is generally okay for me to reply, "He might be," unless I know that he is not. This is just what we should expect if the truth of "He might be" depends on what the speaker knows. It is not what we should expect if the truth of "He might be" depends in part on what others know, or on what one could come to know. As we will see in what follows, the more "objective" we make claims about epistemic possibility, the larger the gap between the circumstances in which we are warranted in making them and the circumstances in which we actually do make them. Solipsistic Contextualism explains why we are willing to assert "It might be that p" in roughly the same cases as "For all I

[1] Solipsistic Contextualism is sometimes attributed to G. E. Moore (perhaps the first philosopher to clearly distinguish epistemic uses of modals from others) on the basis of passages like this one, from his *Commonplace Book*:

> People *in philosophy* say: The *props.* that I'm not sitting down now, that I'm not male, that I'm dead, that I died before the murder of Julius Caesar, that I shall die before 12 to-night, are "logically possible". But it's not English to say, with this meaning: It's possible that I'm not sitting down now etc.—*this* only means "It's not certain that I am" or "I don't know that I am".

However, Moore did not accept the Solipsistic Contextualist analysis of "must." He denied that "It *must* be that p" means the same as "It's impossible that not-p" (188), on the grounds that it is appropriate to say the former only when one does not know *directly* (e.g. by seeing) that p. It seems that he also rejected the solipsistic view for "probably" (402).

know, p."

Second, Solipsistic Contextualism beautifully explains why the following sentences sound paradoxical:

(1) Joe might be in Boston, but I know he isn't.

(2) Joe might be in Boston, but he isn't.

According to Solipsistic Contextualism, (1) is a contradiction: when the second conjunct expresses a truth, the first must express a falsehood. And, while (2) isn't a contradiction—possibility had better not imply actuality!—it is pragmatically infelicitous, since in asserting that Joe isn't in Boston, one represents oneself as knowing that he isn't, contrary to what is conveyed by the first conjunct.[2]

However, there are serious problems with Solipsistic Contextualism. I won't be alone in pointing them out: most of them have been noticed already by nonsolipsistic contextualists and expressivists. But I think that the former have failed to appreciate how deep these problems are, while the latter have appreciated them but overreacted. As I will argue below, once the force of the objections to Solipsistic Contextualism have been properly appreciated, it becomes clear that there is no stable nonsolipsistic fix. Recognizing this, expressivists have abandoned the whole project of doing truth-conditional semantics for epistemic modals. But that is throwing the baby out with the bathwater: there is, as I will argue, a viable truth-conditional semantics for epistemic modals, provided one is willing to entertain the idea that truth varies not just with the context in which a claim is made, but with the context in which it is *assessed*.

2 Against Solipsistic Contextualism

I'll consider three arguments against Solipsistic Contextualism. All of them are facets of a single problem: Solipsistic Contextualism cannot explain why we take ourselves to be disagreeing with each other about what might be the case, even when we have very different bodies of background knowledge.

2.1 Third-person assessments

The first problem is that people don't assess others' epistemic modal claims in the way that they should if Solipsistic Contextualism were correct. They

[2] Cf. DeRose 1991, 600, Stanley 2005.

don't take them to be equivalent to claims about what is ruled out by what the speaker knows at the time of utterance.

I'd like you to imagine yourself in two slightly different scenarios. I'll ask a question about each; write down your answer.

> *First case:* You overhear George and Sally talking in the coffee line. Sally says, "I don't know anything that would rule out Joe's being in Boston right now" (or perhaps, more colloquially, "For all I know, Joe's in Boston"). You think to yourself: *I* know that Joe isn't in Boston, because I just saw him an hour ago here in Berkeley. *Question:* Did Sally speak falsely?
>
> *Second case:* Scene as before. Sally says, "Joe might be in Boston right now." You think to yourself: Joe can't be in Boston; I just saw him an hour ago here in Berkeley. *Question:* Did Sally speak falsely?

Did you answer "No" to the first question and "Yes" to the second? Of course we don't have grounds for supposing that Sally spoke falsely in the first case: she was simply commenting on what she knew. In the second case, though, it seems quite natural to reject her claim as false on the basis of the same information.[3]

Of course, we must take care that we are rejecting Sally's whole claim as false, and not just the (embedded) proposition that Joe *is* in Boston. Compare this dialogue:

> "It's rumored that you are leaving California."
>
> "That's completely false!"

Here the point of the response is to reject the thing that is rumored, not the claim *that* it is rumored. Could something similar be said about our inclination to reject Sally's claim?

[3]This phenomenon was first called to my attention by a footnote in John Hawthorne's book *Knowledge and Lotteries*: "[A]s far as I can tell, ordinary people evaluate present tense claims of epistemic modality as true or false by testing the claim against their own perspective. So, for example suppose Angela doesn't know whether Bill is alive or dead. Angela says *Bill might be dead*. Cornelius knows Bill is alive. There is a tendency for Cornelius to say Angela is wrong. Yet, given Angela's perspective, wasn't it correct to say what she did? After all, when I say *It might be that P and it might be that not P*, knowing that Cornelius knows whether *P*, I do not naturally think that Cornelius knows that I said something false. There is a real puzzle here, I think, but this is not the place to pursue it further." (Hawthorne 2004, 27 n. 68)

We have ways of distinguishing between cases where the whole asserted content is being rejected and cases where the embedded proposition is being rejected. The easiest way is just to ask:

> "Do you mean that it's false that you're leaving California, or that it's false that that's what's rumored?"
>
> "The former."

So, since you are the protagonist in the two cases I described above, let me ask you. When you said (supposing you did) that Sally spoke falsely, did you mean that she spoke falsely in saying "Joe might be in Boston," or just that it's false that Joe *is* in Boston? It was the former, right? Perhaps there would be some ambiguity if you had assented to "That's false." But you assented to "Sally spoke falsely," which clearly concerns what Sally asserted, not its embedded complement.

2.2 Retraction

If that's not enough, try this test: Should Sally *retract* her assertion, or can she stand by it? Consider how odd it would be for your interlocutor in the rumor case to retract her assertion:

> "It's rumored that you are leaving California."
>
> "That's completely false!"
>
> "Okay, then, I was wrong. I take back what I said."

Your interlocutor wasn't wrong about anything and can quite reasonably let her assertion about what is rumored stand:

> "What a relief! But that *was* the rumor."

By contrast, it seems entirely natural for Sally to retract her assertion that Joe might be in Boston after she hears what George has to say:

> "Joe might be in Boston."
>
> "No, he can't be in Boston. I just saw him an hour ago in Berkeley."
>
> "Okay, then, scratch that. I was wrong."

Indeed, it would be very odd for Sally *not* to retract her claim (explicitly or implicitly):

> "Okay, then, he can't be in Boston. But I still stand by what I said a second ago."

It's not plausible to say that the target of Sally's retraction (the thing she takes herself to have been wrong about) is the embedded proposition—that Joe *is* in Boston—for she didn't assert or believe *that*. It must, then, be the modal proposition she expressed by saying "Joe might be in Boston."

It is important here to distinguish *retracting* an assertion from claiming that one ought not to have made it in the first place. To say that one was wrong *in claiming* that p is not to say that one was wrong *to claim* that p. Sometimes it is right to make a claim that turns out to have been wrong (false). For example, suppose that all of the evidence available to Holmes overwhelmingly supports the hypothesis that the butler is the murderer. Then he was not wrong *to* claim that the butler was the murderer, even if it turns out that he was wrong *in* so claiming. Not only was he right to claim that butler was the murderer—following the evidence, as always—but he would have been wrong to withhold his view on the matter.

If you find it implausible that Sally would say "I was wrong" in the dialogue above, make sure you're not interpreting her as saying "I was wrong to say that." Of course she wasn't wrong to say what she did. But what she said was wrong, and that is what she is acknowledging.

2.3 Disputes

Here is a third reason for rejecting Solipsistic Contextualism. It seems that we sometimes *argue* and *disagree* about epistemic modal claims. A conversation might center, for a time, on the question whether Joe might be in Boston. The issue is not whether Joe *is* in Boston; everyone present acknowledges that he might be in Berkeley, and so no one thinks that there are going to be grounds for asserting that he *is* in Boston. The point of the conversation is to settle whether he *might be* in Boston. Reasons are offered on both sides, disputes are resolved, and perhaps a consensus is reached.

It is crucial to such disputes that the participants take themselves to be contradicting each other when one says "It might be that p" and the other says "No, it can't be that p." Solipsistic Contextualism cannot make sense of this. For it holds that the first participant's claim is about what *she* knows, while the second's is about what *he* knows.

A (broadly) Solipsistic Contextualist might account for this data by taking epistemic modals to work the way "local" seems to work. If your brother in Anchorage says "I went to a local bar, the Moose's Tooth," you (in Berke-

ley) can reply: "That's not local, it's five miles away from you!" (meaning *local to your brother*). In the same way, the Contextualist might say, epistemic modals can be used with reference to what someone else (say, one's interlocutor) knows. This move would help make sense of perceived disagreement.

It would do so, however, by construing disputes about what might be the case as disputes about what some particular person knows at some particular time. But then we should expect them to be asymmetrical in a way that they are not, since the person in question has privileged access to what she believes, and this is relevant to what she knows. Disputes about what might be the case do not feel as if they are "centered on" a particular person in this way. Indeed, they feel like continuous arguments, with a single topic, even as the participants gain relevant knowledge through discussion. Solipsistic Contextualism cannot account for this.

2.4 Semantic blindness?

All I am doing here is calling attention to how we use epistemic modals in practice. The defender of Solipsistic Contextualist could always acknowledge these facts but dismiss them as misleading guides to the semantics of epistemic modals. Perhaps third parties who assess Sally's claim mistakenly take her to have asserted what *they* would be asserting by saying "Joe might be in Boston." Perhaps Sally, assessing her own past assertion, mistakenly takes it to have the content she would *now* express if she used the same sentence. And perhaps the parties to a dispute about whether it's possible that Joe is in Boston are mistakenly taking themselves to contradict each other, when in reality they are simply talking past each other.

But that's a lot of error to impute to speakers. One wants some explanation of why speakers are systematically confused in this way, and why this confusion doesn't generalize to other cases that should be similar if Solipsistic Contextualism is correct. For example, if speakers are systematically blind to unobvious context sensitivity, why doesn't the following dialogue seem natural?

> "Joe is tall. In fact, he's the tallest graduate student in our department."
>
> "No, he isn't tall. He's shorter than nearly every NBA player."
>
> "Okay, then, scratch that. I was wrong."

One would also need to explain why the data that seems to support Solipsistic Contextualism (primarily data about when speakers take themselves

to be warranted in making epistemic modal claims) should be taken so seriously, when the data about third-party assessments, retraction, and disputes are just thrown away. There is no clear reason to favor the "positive" data in this way. Quite the contrary, semantics is typically driven more by data about perceived incompatibilities and entailments than by data about when people are willing to accept sentences. I propose, then, to put this approach to defending Solipsistic Contextualism on the back burner, as a last resort should no alternative view prove viable.

3 Nonsolipsistic Contextualism

These problems with Solipsistic Contextualism are relatively well known. Indeed, practically no one who has staked out a serious position on the semantics of epistemic modals defends the view.[4] It is very common, however, to suppose that the problems with Solipsistic Contextualism lie with its solipsism, and that the solution is to move towards a form of contextualism that is less solipsistic and less subjective. If "Joe might be in Boston" doesn't mean "For all I know, Joe is in Boston," perhaps it means "For all *we* know, Joe is in Boston," or "For all we know or could easily come to know, Joe is in Boston." All of these can be thought of as variants on "What is known does not rule out Joe's being in Boston," with different glosses on "what is known."

In this section, I will consider some different ways in which a contextualist might try to meet the objections we have considered by moving away from the strict Solipsistic Contextualist position. I hope to persuade you that these are all bandaids on a gaping wound. The fundamental problem with Solipsistic Contextualism lies with its Contextualism, not its Solipsism.

3.1 Widening the relevant community

According to NONSOLIPSISTIC CONTEXTUALISM, "Joe might be in Boston" expresses a truth just in case what the contextually relevant group knows does not rule out Joe's being in Boston.[5] There are complications about what it means to say that a *group's* knowledge rules something out, but we will skip over these until section 3.3. There are also complications about how these truth conditions can be generated compositionally: these will be discussed further in section 6, but for our purposes here we need not settle them. The

[4] It appears that Stanley 2005, 128 does endorse it.

[5] See e.g. Hacking 1967, 148, Teller 1972, DeRose 1991.

important thing is that we have replaced talk of the speaker's knowledge with talk of the knowledge of a group picked out by features of the context of use (including, on most versions, the speaker's intentions).

Nonsolipsistic Contextualism allows us to make sense of Sally's retraction of her claim in light of George's response, by supposing that the contextually relevant group includes not just Sally but all the parties to the conversation, George included. That would explain why, when Sally learns that George knew things that precluded Joe's being in Boston, she regards her own claim as having been refuted. It would also vindicate George's assessment of Sally's claim as false. Finally, it would make it possible to understand how a group can argue about whether Joe might be in Boston. According to Nonsolipsistic Contextualism, the group is trying to come to a consensus about what its shared knowledge excludes and leaves open.

Moreover, Nonsolipsistic Contextualism can explain the paradoxical ring of sentences (1) and (2) just as well its Solipsistic cousin. For it is usually assumed that the speaker belongs to the contextually relevant group, and that the group counts as knowing if any member does. On these assumptions, if the speaker knows that Joe isn't in Boston, then "Joe might be in Boston" cannot express a truth. It follows that (1) is a contradiction and that (2) is pragmatically infelicitous.

So far, the move away from solipsism seems well-motivated and plausible. The problem is that once we let data about third-party assessments and retraction motivate an expansion of the contextually relevant group to include more than just the speaker, there is no way to stop this machine. The same kind of arguments that motivate expanding the relevant group of knowers to include George (in our example above) will motivate expanding the relevant group of knowers to include anybody who will ever consider the claim.

Indeed, the problem can be seen in our very first example with Sally and George. When you overhear Sally telling George, "Joe might be in Boston," you think to yourself "She has spoken falsely." To make sense of this reaction, the Nonsolipsistic Contextualist will have to make the contextually relevant group of knowers include *you*, even though you are not part of the conversation, not known to Sally, and perhaps not even *noticed* by Sally. It seems, then, that we need to take Sally's claim to concern not just what she and George know, but what anyone within earshot of their conversation knows.

And why limit ourselves to earshot? It doesn't matter much to our story that you are in the same room as Sally. You'd assess her claim the same way if you were thousands of miles away, listening through a wiretap. Indeed, it

seems to me that it does not even matter whether you are listening to the wiretap live or reviewing a recording the next day—or the next year.[6] To vindicate all these third-party assessments, the Nonsolipsistic Contextualist would have to extend the relevant group of knowers not just to those in earshot, but to all those who will one day hear of, read of, or perhaps even conjecture about, Sally's claim. There's no natural stopping point short of that.

Consideration of when speakers will retract their claims seems to point in the same direction. For it seems to me that the retraction data we considered in section 2.2 is just as robust when we replace George by a hidden eavesdropper. Suppose Sally says, "Joe might be in Boston," and George replies, "Oh really? I didn't know that." At this point, Jane—who is hiding in the closet—emerges and says, "Joe can't be in Boston; I just saw him down the hall." It seems entirely natural for Sally to reply, "Oh, then I guess I was wrong. Thanks, Jane." It would be bizarre for her to say, "Thanks for telling us, Jane. I guess Joe can't be in Boston. Nonetheless, I stand by what I said a second ago." Clearly Sally did not have Jane in mind when she made her claim. So if we're going to make sense of these retractions, we must suppose that the force of Sally's claim was something like: *what we know—we who are or will be in a position to consider this claim—does not rule out Joe's being in Boston.*

The same point can be made by considering *disputes* about what might be the case. Suppose two research groups are investigating whether a certain species of snail can be found in Hawaii. Neither group knows of the other's existence. One day they end up at the same bar. The first group overhears members of the second group arguing about whether it is "possible" that the snails exist on the big island, and they join the discussion. Although the two groups have different bodies of evidence, it does not intuitively seem that they are talking past each other when they argue. Nor does it seem as if the topic *changes* when the first group joins the discussion (from what was ruled out by the second group's evidence to what is ruled out by both groups' evidence). To accommodate these intuitions, the Nonsolipsistic Contextualist will have to take *all* the possibility claims made by both groups to concern what is ruled out by the collected evidence of everyone who is investigating the question (known or unknown)—for any of these investigators could show up at the bar, in principle.

To sum up: the arguments that motivate a move from the "for all I know"

[6] In that case it will be your knowledge of Joe's whereabouts on the day the recording was made that is relevant—but still *your* knowledge (not Sally's), and your knowledge *now*.

reading of epistemic modals to the "for all we know" reading also motivate extending the scope of "we" to include not just the participants in the conversation but eavesdroppers, no matter how well hidden or how distantly separated in time and space. "It is possible that p" becomes "p is not ruled out by what is known by anyone who will ever consider this claim."

But this is something like a *reductio ad absurdum* of Nonsolipsistic Contextualism. For if this is what epistemic modals mean, then most ordinary uses of them are completely irresponsible. Surely Sally would not be warranted in asserting "Nothing known by me *or by anyone who will ever consider this claim* excludes Joe's being in Boston." Indeed, she may have good reason to deny this. But intuitively Sally *is* warranted in asserting that Joe might be in Boston; her assertion is a paradigm use of an epistemic modal.

3.2 Objective factors

Hacking 1967 has a somewhat different argument for the same conclusion, that widening the relevant group of knowers to include the speaker's conversational partners will not suffice to save a contextualist semantics for epistemic modals:

> Imagine a salvage crew searching for a ship that sank a long time ago. The mate of the salvage ship works from an old log, makes a mistake in his calculations, and concludes that the wreck may be in a certain bay. It is possible, he says, that the hulk is in these waters. No one knows anything to the contrary. But in fact, as it turns out later, it simply was not possible for the vessel to be in that bay; more careful examination of the log shows that the boat must have gone down at least 30 miles further south. The mate said something false when he said, "It is possible that we shall find the treasure here," but the falsehood did not arise from what anyone actually knew at the time. (148)

Hacking concludes that the truth of epistemic modal claims must depend not just on what is known, but on objective features of the situation—here, the presence of relevant information in the log.

This is another way in which contextualism might be made nonsolipsistic: instead of (or in addition to) widening the community of relevant epistemic agents, we relax the strength of the relation these agents must stand in to the relevant facts. In addition to looking at what they *do* know, we look at what they could *come* to know through a "practicable investigation" (as Hacking puts it), or what is within their "epistemic reach" (as Egan 2007

puts it). We might say that "it is possible that p" expresses a truth if what is within the speaker's epistemic reach (or perhaps the epistemic reach of a contextually relevant group) does not rule out p. Similar ideas can be found in DeRose 1991, which talks of "relevant way[s] by which members of the relevant community can come to know," and even in G. E Moore's *Commonplace Book*.[7]

On this view, the reason Sally speaks falsely when she says "Joe might be in Boston" is that she has within her "epistemic reach" facts that would have ruled out Joe's being in Boston. A "practicable investigation"—simply asking those around her—would have settled the matter. That also explains why Sally retracts her assertion when she hears what George has to say. Finally, it explains how it is that a group of people can argue about "whether Joe might be in Boston" without talking past each other or constantly changing the subject as they learn new things. The real topic is whether the facts that are within the group's "epistemic reach" suffice to rule out Joe's being in Boston.

I am skeptical that speakers make any implicit distinction in their use of epistemic modals between "practicable" and "impracticable" investigations, or between what they can easily come to know and what they can come to know only with difficulty or by the cooperation of fate. For example, it seems correct to say that people who used to think that it was possible that there were even numbers greater than 2 and less than 10^{17} that were not the sum of two primes were wrong—since we have now verified computationally that there cannot be any such numbers—even though this computation was not a practicable investigation for *them*. Similarly, we will judge Sally's claim false (on the basis of what we know) even if we are listening in remotely, so that Sally is unable to take advantage of our information about Joe's wherabouts. And Sally will retract her assertion that Joe might be in Boston just as surely if she finds an itinerary on the floor as she will in response to George's intervention—even if her finding this scrap of paper is completely fortuitous and not the result of a "practicable investigation" or a contextually relevant "way of coming to know."

Even leaving this worry aside, however, it seems to me that Hacking's is the wrong fix. Consider his own salvage ship example. It seems perfectly reasonable for the mate to say:

[7]Moore writes: "Things which no-one in fact knows may be such that, owing to them, it is in fact likely or unlikely that p, provided they are such that the person who says p is likely or unlikely *easily might* know, or which the speaker & his hearers *couldn't easily know or have known,* is incompatible with p, doesn't prevent its being true that p is prob." (Moore 1962, 402, emphasis added).

> "It's possible that we shall find the treasure here, and it's possible that we shall find it farther south. Let's examine the log before we dive: maybe we can eliminate one of these locations."

In his second sentence, the mate is acknowledging the possibility that a "practicable investigation" will rule out one of the two possibilities. If Hacking is right, that is tantamount to acknowledging that one of the two conjuncts of the mate's first sentence might be false. So if Hacking's proposal is right, then the mate's speech should sound as infelicitous as "Jane is in Boston and Al is in New York. Maybe Jane is not in Boston." But it doesn't; it is perfectly felicitous.

3.3 Distributed knowledge

A different way in which one might handle cases like Hacking's, in which an epistemic modal claim seems to be false even though the proposition said to be possible is not ruled out by what anyone knows, is to appeal to *distributed* knowledge. We have been appealing, vaguely, to "what is known by a contextually relevant group G." But what is it for a group G to know that p? A variety of answers are possible:

> *Universal knowledge:* Every member of G knows that p.
>
> *Partial knowledge:* Some member of G knows that p.
>
> *Common knowledge:* Every member of G knows that p, and knows that the other members know that p, and that they know that the other members know that p, etc.
>
> *Distributed knowledge:* p is a consequence of the totality of facts known by various members of G.

Teller 1972 suggests that if we take epistemic modal claims to concern a group's *distributed* knowledge, we can explain why claims of the form "It is possible that p" sometimes seem false even though no one in the speaker's group is in a position to rule p out.

Consider the unfortunate murder of McRich (Teller 1972, 310). Sleuth knows that McRich's nephew was ten miles from the scene of the crime all evening, while Private Eye knows that the murder occured between 7 and 8 p.m. Both believe that it's possible that the nephew did it. When they compare notes, they realize that the nephew couldn't be the murderer. Teller points out how natural it would be for them to concede that they were wrong before, and that it had only *seemed* possible that the nephew was the murderer. The explanation, on Teller's view, is that the truth of their claims of

epistemic possibility depends on what is known *distributively* by the two of them together, which rules out the possibility that the nephew is the murderer. "What we know," in this sense, can include facts not known to any of us individually.

Like broadening "epistemic reach," appealing to distributed knowledge in the semantics for epistemic modals can make epistemic modal claims more "objective." This helps account for the fact that we tend to assess them in light of information not possessed by the speaker or any members of the speaker's group. The problem, as before, is that it threatens to make them *too* objective. Given that Sleuth and Private Eye both have reason to believe that the other has information he does not have, it would be rash for either to assert or believe that what is known *distributively* by them fails to rule out the nephew as murderer. So if Teller is right about epistemic modals, it should seem rash for either of them to assert or believe that it's possible that the nephew did it. But it doesn't seem rash. It seems perfectly appropriate.

3.4 The puzzle

All of the proposals we've considered in this section are attempts to keep the core contextualist idea of Solipsistic Contextualism—the idea that epistemic modals are contextually sensitive to what is known at the context of use—while dropping the implausible Solipsism. And all of them face the same basic problem. The less solipsistic the theory becomes, the harder it is to explain why speakers feel entitled to make the epistemic modal claims they do.

The problem is that we have two kinds of data, and they seem to point in different directions. If we attend to facts about when speakers take themselves to be warranted in asserting that something is "possible," Solipsistic Contextualism looks like the right view. Unfortunately, it cannot account for the data about speakers' assessments of epistemic modal claims—including self-assessments that prompt retraction—or for the nature of disputes about questions expressed using epistemic modals. We can account for these data by making our Contextualism less solipsistic, but then we can no longer account for the data that originally motivated Solipsistic Contextualism.

Nor does there seem to be any stable position that balances these two competing desiderata. If we focus on *uptake* (third-party assessments, retractions, and disagreement), we are led to expand the relevant body of knowledge, seemingly without end. But if we focus on *production*, we are led to contract it (on pain of making ordinary, apparently reasonable assertions unwarranted). We are led to a kind of paradox: although the truth

of a claim made using epistemic modals must depend somehow on what is known—that is what makes it "epistemic"—it does not seem to depend on any *particular* body of knowledge. And there is no way to account for this in the framework of contextualism, which requires that the relevant body of knowledge be determined by features of the context of use. The fundamental problem with Solipsistic Contextualism lies with its Contextualism, not its Solipsism.

4 Non-truth-conditional Approaches

If these arguments seem familiar, perhaps it's because they've been made before. Consider how Price 1983 argues against truth-conditional treatments of "probably." First, he points out that we do not treat claims about what is "probable" as claims about what is likely given the *speaker's* evidence:

> If I disagree with your claim that it is probably going to snow, I am not disagreeing that given *your* evidence it is likely that this is so; but indicating what follows from *my* evidence. Indeed, I might *agree* that it is probably going to snow and yet think it false that this follows from your evidence. (403)

He then notes that if we fix this problem by expanding the relevant body of evidence to include, say, evidence that is available in principle, we can no longer understand how speakers take themselves to be justified in making the probability judgements they do:

> . . . consider the surgeon who says, 'Your operation has probably been successful. We could find out for sure, but since the tests are painful and expensive, it is best to avoid them.' The accessibility, in principle, of evidence which would override that on which the SP judgement is based, is here explicitly acknowledged. (405)

If we look at when speakers make "probably" claims, we are pushed towards a solipsistic semantics, while if we look at third-party assessments of such claims, we are pushed toward something more objective. The upshot is that there is no way of filling in the X in "Given evidence X, it is probable that q" that would yield plausible truth conditions for the unqualified "It is probable that q."

Price takes these arguments to be compelling reasons for the view that "probably" does not contribute to the propositional content of a speech

act at all. His view is that "probably" contributes to the *force* of a speech act, not its content.[8] Other philosophers and linguists have taken similar views about "possibly" and other epistemic modals. So it is worth considering whether such approaches might provide a satisfactory resolution to the problems scouted in the preceding two sections.

4.1 Epistemic modals as force modifiers

It would be misguided to ask how "speaking frankly" contributes to the truth conditions of

(3) Speaking frankly, she's too good for him.

When (3) is used to make an assertion, what is asserted is simply *that she's too good for him*. "Speaking frankly" does not contribute anything to the content of the assertion; its role is rather to comment on the kind of speech act being made. We should not puzzle ourselves about when the proposition *that speaking frankly she's too good for him* is true, because there is no such proposition.

Perhaps asking how epistemic modals affect truth conditions is equally misguided. We have assumed so far that Sally is making an assertion, and this assumption leads directly to questions about the truth conditions of her claim. But we need not understand her speech act as an assertion. Perhaps she is simply signalling her unwillingness to assert that Joe *isn't* in Boston. As Hare argues, "We have a use for a way of volubly and loquaciously *not* making a certain statement; and perhaps there is one sense of 'may' in which it fulfils this function" (1967, 321). Or perhaps she is *perhapserting* the proposition *that Joe is in Boston*. Here a "perhapsertion" is a distinct kind of speech act, which we might understand as the expression of some minimal degree of credence, or advice not to ignore a possibility. If the linguistic role of epistemic modals is to signal that the speaker is making a perhapsertion, then we need not trouble ourselves about the contribution it makes to truth conditions.

Such views account quite well for our uses of (standalone) sentences involving epistemic modals, while allowing us to dodge the questions about the truth-conditional contribution of epistemic modals that we saw above to be so problematic. However, they leave us unequipped to deal with *embedded* uses of epistemic modals. And in general, they make it difficult to

[8] In later work (1994) he suggests that the speech act can be *both* an assertion that it might be that p—in some minimal sense of "assertion"—*and* a non-assertive expression of positive credence in p.

explain interactions between epistemic modals and expressions that have a content-expressing role.

4.2 Interface problems

Epistemic modals can occur embedded under quantifiers, truth-functional connectives, conditionals, attitude verbs, adjectives, and other constructions.[9] In this they differ greatly from "speaking frankly," which does not embed in these ways:

(4) (a) If it might be raining, we should bring umbrellas.
(b) #If speaking frankly she's too good for him, she'll realize this.

(5) (a) It's not possible that Joe is in Boston.
(b) #It's not the case that speaking frankly, Joe is in Boston.

(6) (a) Sally believes that it's possible that Joe is in Boston.
(b) #Sally believes that speaking frankly, she's too good for him.

The force modifier approach tells us nothing about the contribution made by "might" in (4a) or "possible" in (5a). It is clear that "might" in (4a) is *not* indicating that anything is being perhapserted. In typical uses of (4a), the whole conditional is being asserted full stop, and the antecedent is neither asserted nor perhapserted. (It's perfectly coherent to say, "If p, then q. But not p.") There is clearly a difference between (4a) and

(7) If it is raining, we should bring umbrellas,

but the force-modifier account of "might" does not help us understand what it is, since "might" is not serving as a force modifier in (4a).

Similarly, the force-modifier account of

(8) It's possible that Joe is in Boston

gives us no guidance whatsoever about the meaning of (5a). Clearly "possible" occurs here within the scope of the negation—(5a) does not mean the same thing as

(9) It's possible that Joe is not in Boston

[9]However, there are some interesting restrictions. For example, von Fintel and Iatridou 2003 argue that in many contexts epistemic modals must take wide scope over quantifiers.

—but what sense can we make of the negation of a *speech act*?

Finally, in (6a), "possible" occurs in the description of the content of a cognitive state, not a speech act. Although it is fairly clear how we could leverage our understanding of the kind of speech act conventionally made by (8) into an understanding of (6a), this requires that we treat "believe" differently when its complement is modified by an epistemic modal than when it is not. (Roughly: when "believes" takes a complement clause in which an epistemic modal takes wide scope, it will attribute credence above some minimal threshold, while in other cases it will attribute full belief.) Similar modifications will be needed for other attitude verbs. This complicates the (already difficult) project of giving a compositional semantics for attitude verbs by undermining the neat division of labor between force (supplied by the attitude verb) and content (supplied by the complement clause).

An advocate of the force-modifier approach might be able to tell separate stories, like the story sketched above about attitude verbs, about how epistemic modals behave in all of these other embedded contexts. But the resulting account is bound to be ugly and complex. The beauty of truth-conditional semantics is that it provides a common currency that can be used to explain indefinitely many interaction effects in a simple and economical account. We should be prepared to accept a messy, non-truth-conditional account of epistemic modals only if there is no truth-conditional account that explains the data.

4.3 Explaining retractions

In addition to these problems with embedded uses, the force-modifier approach has difficulty with the same retraction data that caused problems for contextualism. For, if the force-modifier view is right, why does Sally say "I was wrong" when George tells her about Joe's whereabouts? None of the answers that are available on the force-modifier view seem to work:

1. *She believed that Joe was in Boston, and he wasn't.* No, because she didn't believe this.

2. *She had a minimal degree of credence that Joe was in Boston, and he wasn't.* No, because there's nothing "wrong" about having a minimal degree of credence in a proposition that turns out to be false. For example, it's quite reasonable to have a minimal degree of credence in each of a number of incompatible alternatives, even though all but one of these are bound to be false.

3. *She had a minimal degree of credence that Joe was in Boston, and she shouldn't have, given her evidence.* But she should have! *Her* evidence didn't rule out his being in Boston.

4. *She raised to salience the possibility that Joe was in Boston, and she shouldn't have.* But she should have! It was reasonable and appropriate for her to do so.

In order to exhibit Sally's retraction as rational, we need to understand how she can reasonably take herself to have performed a speech act that is in some way incorrect. The force-modifier approach lacks the resources to do this.

5 A "relativist" approach

Advocates of force-modifier accounts are typically well aware of the interface problems canvassed in the last section. That is why they motivate their views by arguing *against* truth-conditional approaches. For example, Simon Blackburn says that although his expressivist theory of evaluative language will no doubt have "Ptolemaic" complexities, there is no "Copernican" theory that explains the data better (Blackburn 1984, 195–6). Price's argument for a force-modifier approach to "probably" proceeds along similar lines.

Such arguments work only if they can rule out all possible truth-conditional approaches. Typically, they assume that any such truth-conditional view must have a contextualist shape. In the case of epistemic modals, this means that the body of known facts relative to which the modal is assessed must be determined by features of the context of use (including the speaker's intentions). We have seen above how one might argue quite generally that no view with this shape accurately captures the way we use epistemic modals.

But must a truth-conditional semantics for epistemic modals have this shape? In this section, I want to explore the possibility of broadening our semantic frameworks to make room for a new kind of view, on which the truth of epistemic modal claims depends on a body of known facts determined not by the context of use, but by what I'll call the *context of assessment*. This semantics offers prospects for meeting the objections to contextualist views in a broadly truth-conditional framework, thereby undermining the motivation for the force-modifier approach.

5.1 Bicontextuality

We can understand the notion of a context of assessment by analogy with the familiar notion of a context of use:

Context of use: the setting for an actual or possible use of a sentence (or proposition) in a speech act or mental act.

Context of assessment: the setting from which such a use is being assessed for truth or falsity on some actual or possible occasion of assessment.

For many purposes, one can think of a context as a centered possible world—a world-time-agent triple—since all of the other contextual factors that are needed are determined once a centered world is given. We can then talk of "the speaker of the context of use," "the time of the context of assessment," or "the epistemic state of (the assessor at) the context of assessment." Alternatively, one can think of a context of assessment as an abstract sequence of parameters representing semantically relevant features of a (concrete) setting from which a speech act or other use of a sentence might be assessed. I will take the first approach here (following Lewis 1980 rather than Kaplan 1989), but nothing hangs on it.

Since we do assess uses of sentences, and whenever we do this we occupy some particular context, there is little to object to in the concept of a "context of assessment." Semanticists of all stripes should be able to deploy this concept; the only question is whether it has a useful role to play. The question is whether the truth, reference and other semantic properties can depend not just on features of the context in which a sentence is used, but on features of the context in which it is assessed. To answer Yes to this question is to acknowledge a new kind of context sensitivity, which I have called *assessment sensitivity* to distinguish it from the familiar *use sensitivity*.[10]

It should be obvious where this is going. We started with the intuitively compelling idea that the truth of epistemic modal claims depends on *what is known*. That is why they are called "epistemic." But we ran into trouble when we tried to answer the question, "known to whom?" For it seemed that people tend to assess epistemic modal claims for truth in light of what *they* (the assessors) know, even if they realize that they know more than the speaker (or relevant group) did at the time of utterance. A straightforward way to account for this puzzling fact is to suppose that epistemic modals

[10] A sentence (or proposition) is *use-sensitive* iff its truth as used at c_U and assessed at c_A depends on features of c_U. A sentence (or proposition) is *assessment-sensitive* iff its truth as used at c_U and assessed at c_A depends on features of c_A.

are assessment-sensitive: the truth of an epistemic modal claim depends on what is known by the assessor, and thus varies with the context of assessment. On this view, epistemic modal claims have no "absolute" truth values, only assessment-relative truth values. This is why they resist being captured in standard frameworks for truth-conditional semantics.

For the sake of concreteness, we'll work at first with the most austere kind of relativist view—what one might call SOLIPSISTIC RELATIVISM. (Later we'll consider some complications.) On this view, "Joe might be running" expresses a truth, as assessed by Sam, just in case what *Sam* knows (at the time of assessment) does not rule out that Joe is running. This is not yet a compositional semantics for "might," since we have not explained how to handle embedded occurrences. More on that later (section 6). But we can already see from this sketch of a theory how Solipsistic Relativism will handle the data that seemed most problematic for the various forms of contextualism.

5.2 Explaining third-party assessments

Solipsistic Relativism has a very straightforward explanation of the data about third-party assessments. According to Solipsistic Relativism, the truth of an epistemic modal claim (relative to a context of assessment) depends on what the assessor knows, not what the speaker knew when making the claim. So it is appropriate for eavesdroppers to assess the truth of epistemic modal claims against the background of what they know, even if this is very different from what the speaker knew.

Recall that the contextualist could only handle the eavesdropper data by strengthening truth conditions for claims of epistemic possibility to the point where it became hard to understand why people would make them at all. The relativist does not have this problem. Sally's claim that Joe might be in Boston is true as assessed from the context in which she makes it, so we can understand why she makes it in the first place. In general, Solipsistic Relativism counts a sentence as true as used at c and assessed at c just when Solipsistic Contextualism counts it as true as used at c. The relativist semantics will diverge from the contextualist semantics only when the context of assessment is distinct from the context of use. So the Solipsistic Relativist will be able to explain *production* of epistemic modals in much the same way as the Solipsistic Contextualist, while explaining *assessments* in a way that is not available to the contextualist.[11]

[11] This needs some qualification, since it's not clear that deliberation about whether to assert an assessment-sensitive proposition shouldn't take into account its truth value relative to contexts of assessment other than the one occupied by the speaker. For example,

Hacking's salvage ship case can be handled in the same way. It is really just another third-party assessment case, in which *we* (Hacking's readers) are the third party. According to Solipsistic Relativism, the truth of the mate's claim (as assessed by us) depends on what *we* know. Since we know (from Hacking's narrative) that the treasure lies elsewhere, the mate's claim is false, relative to the context of assessment we occupy. That explains quite straightforwardly why we judge it to be false. The fact that there was a "practicable investigation" the mate could have carried out is simply irrelevant. What is crucial is something Hacking did not explicitly point out: that we, the readers, come to know, through Hacking's testimony, that the treasure lies elsewhere.

5.3 Explaining retractions

The Solipsistic Relativist has an equally simple explanation of why Sally should retract her claim in response to George's correction (section 2.2, above). After Sally learns from George that Joe is not in Boston, she occupies a context of assessment relative to which her original claim is false (since she now knows more than she did). So it is proper for her to retract it.[12]

Note the change of perspective. The contextualist assumes that if what George says implies that Sally's claim is false, then George must be part of the group whose knowledge matters to the truth of Sally's claim. But as we have seen, this way leads to madness: there is no way to keep the group from expanding indefinitely. The relativist, by contrast, holds that what is important is not that *George* knew that Joe was in Berkeley, but that *Sally* comes to know this.

Hence it is irrelevant, for the relativist, that Sally comes to know this through the testimony of someone else who already knew it (at the time she made the claim). What is known by others is relevant only insofar as they are potential informants of the speaker (in this case, Sally). If they don't speak up, or if they do speak up but Sally doesn't believe them (and so doesn't acquire knowledge), then Sally has no objective reason to retract her assertion. Conversely, if the way Sally comes to know something incompatible with Joe's having been in Boston is not through others' testimony but through

one might refrain from asserting something one knows one will have to retract almost immediately, when one's context changes, even if it is true relative to one's current context.

[12] Here I am relying on the normative account of assertion developed in MacFarlane 2005b, according to which one is obligated to retract an assertion that has been shown to have been false, relative to one's current context of assessment.

her own observation, or through serendipitous discovery of evidence, she has just as much reason to retract her original claim, and it seems just as natural for her to do so. That the contextualist isn't getting the right generalization here comes out clearly in the need for epicycles: for example, the appeal to "contextually relevant ways of coming to know" and "distributed knowledge" in addition to a "contextually relevant group of knowers."

5.4 Explaining disputes

As we have seen, the contextualist has difficulty accounting for the fact that people take themselves to be arguing and disagreeing about epistemic modal claims. What are they arguing about? Not about what some particular one of them knows. Perhaps, then, what the group knows. But what if another group joins the discussion? This should seem like a change of subject, and it doesn't. Their disputes seem to concern a common topic—say, whether it is possible that infected birds have entered Alameda county—and this topic can't be reduced to a question about what anyone, or any group, knows.

The Solipsistic Relativist gets this right. On the relativist's account, epistemic modal claims aren't equivalent to any claims about what people know. The former are assessment-sensitive, and the latter are not.[13] The relativist can say that every group that is debating whether it is possible that infected birds have entered Alameda county (by such and such a date) is debating the truth of the same proposition. It's just that the truth of this proposition is perspectival.

5.5 Philosophical debts

Let's take stock. In sections 2 and 3, we saw that contextualist semantics is structurally unable to explain our use of epistemic modals. In order to explain third-party assessments, retraction, and disputes, we need to widen the contextually relevant group of knowers—perhaps indefinitely—and put further, "objective" conditions on the truth of epistemic modal claims. But when we do this, it becomes impossible to explain our readiness to make epistemic modal claims even in situations where we are well aware that others may know more than we do. Historically, this problem was one moti-

[13] Not in the same way, anyway. In MacFarlane 2005a, I argue that knowledge-attributing sentences are assessment-sensitive, because their truth (relative to a context of assessment) depends on the assessor's epistemic standards. But even if this is right, their truth is not sensitive to the same features of contexts of assessment as epistemic modals, so they still won't be equivalent to any epistemic modal claims.

vation for the view that epistemic modals should be understood non-truth-conditionally, as modifiers of the force of a speech act rather than its content. However, as we saw in section 4, this project requires a piecemeal account of the role of epistemic modals in embedded contexts. Such an account, if possible at all, is likely to be very complex. Moreover, force-modifier accounts don't do any better than Solipsistic Contextualism in explaining the retraction data. So we are left with no good account of the meanings of epistemic modals.

Solipsistic Relativism offers a way out. It neatly explains the data that proved impossible to accommodate in a contextualist framework, and it does so without giving up the advantages of a truth-conditional framework. But is it intelligible? If we are to use an assessment-relative truth predicate in our semantic theories, we must pay some philosophical debts. At the very least, we must answer these questions:

1. What changes does relativism require in standard theories of propositions, standard accounts of assertion and belief, and standard approaches to compositional semantics?

2. Isn't this kind of relativism about truth self-undermining, for reasons given by Plato in the *Theaetetus* and repeated by many philosophers since?

3. Even if talk of truth relative to a context of assessment is not self-undermining, do we really understand it? What is it to commit oneself to the truth of an assessment-sensitive proposition? Can the relativist make sense of the idea that belief "aims at" truth?

4. Can we really make sense of *disagreement* about assessment-sensitive claims? If so, what is the point of disagreeing about things whose truth is relative?

5. More broadly, what purpose is served by assessment sensitivity? What would we be lacking if we replaced our assessment-sensitive expressions with assessment-invariant ones (not talking about what might be the case, for example, but only about what various people do and do not know about it)?

I will not try to answer these questions here. I have addressed the first three in MacFarlane 2005b and the last two in MacFarlane 2007. But there is much more clarificatory work to be done before we can be confident that we understand what we are saying when we characterize a claim as true "relative to a context of assessment."

6 Compositional Semantics

The rough characterization of Solipsistic Contextualism in the previous section refers only to standalone sentences in which the epistemic modal takes widest scope. But of course epistemic modals can also occur embedded under quantifiers, conditionals, and other kinds of operators. Since one of the advertised advantages of relativist semantics over the force-modifier approach is its capacity to explain embedded uses, it's worth looking at how standard semantic frameworks must be modified in order to make room for assessment sensitivity, and how a compositional semantics for epistemic modals might look in such a framework.

6.1 Baseline: solipsistic contextualism

As a baseline for comparison, let's start with a version of Solipsistic Contextualism. The aim is to give a finite definition of "true at context of use c" for a first-order language containing the epistemic modal operator "*Might* : " ("it is possible that") and an operator "FAK_t^x :" ("for all x knows at t"). Since "*Might* : ", "FAK_t^x :", and the quantifiers are not truth-functional, we can't simply give a recursive definition of truth at a context of use. Instead, we'll give a recursive definition of truth at a *point of evaluation*, then define truth at a context of use in terms of truth at a point of evaluation.[14] Here a *point of evalution* is an ordered quadruple $\langle c, w, i, a\rangle$, where c is a context, w a possible world,[15] i a set of possible worlds representing an *information state* (intuitively, the worlds not ruled out by the information), and a an *assignment* of objects from the domain relevant at c to the variables.

First, we define the extensions of the primitive terms and predicates of the language, relative to a point of evaluation:

- The extension of "Joe" at $\langle c, w, i, a\rangle$ = Joe.
- The extension of "I" at $\langle c, w, i, a\rangle$ = the agent of c.
- The extension of "now" at $\langle c, w, i, a\rangle$ = the time of c.
- The extension of "human" at $\langle c, w, i, a\rangle$ = the set of humans in w.
- And so on (finitely many of these).

[14] Cf. Tarski 1944, §11, Kaplan 1989, 547.

[15] I won't worry here about how these worlds are to be individuated or whether the same set of worlds can be used in semantics for alethic modals. Though these are important questions, they cross-cut the questions of primary concern to us here.

We can now define truth at a point of evaluation recursively as follows:

- $\ulcorner \phi(\alpha) \urcorner$ is true at $\langle c, w, i, a \rangle$ iff the extension of α at $\langle c, w, i, a \rangle$ belongs to the extension of ϕ at $\langle c, w, i, a \rangle$. (And similarly for polyadic predicates.)
- $\ulcorner \neg \Phi \urcorner$ is true at $\langle c, w, i, a \rangle$ iff Φ is not true at $\langle c, w, i, a \rangle$.
- $\ulcorner \Phi \wedge \Psi \urcorner$ is true at $\langle c, w, i, a \rangle$ iff Φ is true at $\langle c, w, i, a \rangle$ and Ψ is true at $\langle c, w, i, a \rangle$.
- $\ulcorner \exists \alpha \Phi \urcorner$ is true at $\langle c, w, i, a \rangle$ iff for some assignment a' that agrees with a on every variable except possibly α, Φ is true at $\langle c, w, i, a' \rangle$.
- $\ulcorner \mathit{Might} : \Phi \urcorner$ is true at $\langle c, w, i, a \rangle$ iff for some w' in i, Φ is true at $\langle c, w', i, a \rangle$.[16]
- $\ulcorner \mathit{FAK}^{\alpha}_{\tau} : \Phi \urcorner$ is true at $\langle c, w, i, a \rangle$ iff Φ is true at $\langle c, w', i', a \rangle$, where i' is the set of worlds not excluded by what is known by the extension of α at $\langle c, w, i, a \rangle$ at w and the time denoted by τ at $\langle c, w, i, a \rangle$, and w' is some world in i'.

Finally, we can define truth at a context in terms of truth at a point of evaluation:[17]

> An occurrence of a sentence Φ at a context c is true iff Φ is true at every point of evaluation $\langle c, w_c, i_c, a \rangle$, where
>
> - w_c = the world of c,[18]
> - i_c = the set of worlds that aren't excluded by what is known (at c) by the agent of c,
> - a = an assignment of objects from the domain relevant at c to the variables.

[16] Here I depart from the more standard approach of quantifying over the domain of worlds w' such that Rww', where R is a contextually determined "accessibility relation." I have chosen to formulate Solipsistic Contextualism this way, with the domain of worlds provided by a separate information state parameter, in order to make Solipsistic Contextualism easier to compare with Solipsistic Relativism, which uses such a parameter (for the essential difference between the views, see section 6.3, below). There are, however, some excellent reasons for doing the semantics this way, even if one does not want to be a relativist in the end (see Yalcin 2007).

[17] Compare Kaplan 1989, 522. In what follows, I'll use "Φ is true at c" interchangeably with "an occurrence of Φ at c is true."

[18] The assumption that there is a unique "world of c" might prove problematic on some ways of thinking of the epistemic "worlds." I'm not going to pursue this issue further here.

Note that truth at a point of evaluation is defined for all formulas, but truth at a context of use is defined only for sentences (formulas with no free variables).

Let's verify that this account accords with the rough initial statement of Solipsistic Contextualism from section 1. Let Φ be the sentence "*Might* : *Joe is running*," let c be a context in which George is uttering Φ, and let i_c be the set of worlds left open by what George knows at c. Our definition of truth at a point of evaluation tells us that Φ is true at a point of evaluation $\langle c, w, i, a\rangle$ just in case there is some world $w' \in i$ such that "Joe is running" is true at $\langle c, w', i, a\rangle$. Feeding this into our definition of truth at a context, we get that an occurrence of Φ at c is true just in case there is some world w' in i_c such that "Joe is running" is true at $\langle c, w', i_c, a\rangle$ for all assignments a. In other worlds, just in case what George knows at c does not rule out the truth of "Joe is running."

Logical truth and logical consequence can be defined (after Kaplan) as truth and truth preservation at every context:

> A sentence Φ is *logically true* iff for every possible context of use c, Φ is true at c.
>
> A sentence Φ is a *logical consequence* of a set Γ of sentences iff for every possible context of use c, if every member of Γ is true at c, then Φ is true at c.

It is also useful to define a notion of logical necessity that quantifies over points of evaluation rather than contexts, and a corresponding notion of logical implication:[19]

> A formula Φ is *logically necessary* iff for every point of evaluation π, Φ is true at π.
>
> A formula Φ is *logically implied* by a set Γ of formulas iff for every point of evaluation π, if every member of Γ is true at π, then Φ is true at π.

If a sentence is logically necessary, it is logically true, but the converse is not guaranteed. Similarly, if Φ is logically implied by Γ, it is a logical consequence of Γ, but the converse is not guaranteed.

[19] On the need for these two distinct notions, see Thomason 1970, 273 and Kaplan 1989, 548–50. Note that logical truth and consequence are defined only for sentences (closed formulas), while logical necessity and implication are defined for (open or closed) formulas.

Using these definitions, we can show that $\ulcorner FAK^I_{now} : \Phi \urcorner$ and $\ulcorner Might : \Phi \urcorner$ are equivalent in the sense that each is a logical consequence of the other.[20] This is a nice result, because Solipsistic Contextualism was motivated in large part by the intuition that "It might be that p" and "For all I know, p" are in some strong sense equivalent. (Note that they are not equivalent in the stronger sense of logically *implying* each other, for there are points of evaluation at which one is true and the other false. To see this, note that the truth value of $\ulcorner FAK^I_{now} : \Phi \urcorner$ at a point of evaluation $\langle c, w, i, a \rangle$ does not depend at all on the value of i, while the truth value of $\ulcorner Might : \Phi \urcorner$ at that point does depend on the value of i. This makes a difference in embedded contexts: for example, "For all John knows now, for all I know now it is raining" can diverge in truth value from "For all John knows now, it might be raining." $\ulcorner FAK^I_{now} : \Phi \urcorner$ is, however, *strongly* equivalent to $\ulcorner FAK^I_{now} : Might : \Phi \urcorner$: they are true at just the same points of evaluation. This, too, is satisfying, insofar as we seem to use these forms interchangeably in English.)

6.2 Nonsolipsistic contextualism

If we want to make our semantics less solipsistic, it's very easy to do. We can leave everything in the recursive definition of truth at a point of evaluation just as it is. All we need to change is the definition of truth at a context:

> An occurrence of a sentence Φ at a context c is true iff Φ is true at every point of evaluation $\langle c, w_c, i_c, a \rangle$, where
>
> - w_c = the world of c,
> - i_c = the set of worlds that aren't excluded by what is known *distributively* at c by *the group of knowers relevant at c,*
> - a = an assignment of objects from the domain relevant at c to the variables.

[20] Proof: Let c be any context. Let w_c be the world of c, t_c the time of c, s_c the agent of c, and i_c the set of worlds not excluded by what s_c knows at c. Let a be an arbitrary assignment: since we won't be dealing with open formulas, any formula that is satisfied by a can be assumed to be satisfied by any assignment. By the definition of truth at a context, $\ulcorner FAK^I_{now} : \Phi \urcorner$ is true at c iff it is true at the point $\langle c, w_c, i_c, a \rangle$. By the recursive clause for *FAK* :, $\ulcorner FAK^I_{now} : \Phi \urcorner$ is true at $\langle c, w_c, i_c, a \rangle$ iff Φ is true at $\langle c, w', i', a \rangle$, where i' is the set of worlds not excluded by what is known by the extension of "I" at $\langle c, w_c, i_c, a \rangle$ at w_c and the time denoted by "now" at $\langle c, w_c, i_c, a \rangle$, and w' is some world in i'. But the extension of "I" at $\langle c, w_c, i_c, a \rangle$ is s_c and the time denoted by "now" at $\langle c, w_c, i_c, a \rangle$ is t_c. So $i' = i_c$. Thus $\ulcorner FAK^I_{now} : \Phi \urcorner$ is true at c iff for some world $w' \in i_c$, Φ is true at $\langle c, i_c, w', a \rangle$. But as we have seen, this is just the condition for $\ulcorner Might : \Phi \urcorner$ to be true at c.

We could add "objective factors" just as easily—again, by modifying the clause for i_c in the above definition. Or we could refrain from specifying exactly how the relevant information state is determined by features of context, and say simply:

- i_c = the information state relevant at c.

The differences between the various sorts of contextualism would then be cast as differences about what makes an information state "relevant" at a context. We might call this formulation FLEXIBLE CONTEXTUALISM.

6.3 Solipsistic relativism

Moving to a view on which "might" is assessment-sensitive is nearly as easy. Again, we need not modify the recursive definition of truth at a point of evaluation. The only change needed is in the definition of truth at a context—or, now, at context*s*, for we can only ask about the truth of an occurrence of a sentence relative to some particular context of assessment. To move from Solipsistic Contextualism to Solipsistic Relativism, we need only substitute the context of assessment for the context of use in the clause governing the initialization of the i parameter:

> An occurrence of a sentence Φ at a context c_U is true as assessed from a context c_A [21] iff Φ is true at every point of evaluation $\langle c_U, w_{c_U}, i_{c_A}, a\rangle$, where
>
> - w_{c_U} = the world of c_U,
> - i_{c_A} = the set of worlds that aren't excluded by what is known (at c_A) by the agent of c_A,
> - a = an assignment of objects from the domain relevant at c to the variables.

It is trivial to verify that these definitions yield the result described earlier, that $\ulcorner$*Might* : $\Phi\urcorner$ is true as used at c_U and assessed at c_A iff what is known to the assessor at c_A is compatible with the truth of Φ at $\langle c_U, c_A\rangle$.

Logical truth and consequence can be defined as before, only we quantify over both contexts of use and contexts of assessment:

[21] In what follows, I'll use "i is true at context of use c_U and context of assessment c_A" interchangeably with "an occurrence of i at c_U is true as assessed from c_A."

> A sentence Φ is *logically true* iff for every possible context of use c_U and context of assessment c_A, Φ is true as used at c_U and assessed from c_A.
>
> A sentence Φ is a *logical consequence* of a set Γ of sentences iff for every possible context of use c_U and context of assessment c_A, if every member of Γ is true as used at c_U and assessed from c_A, then Φ is true as used at c_U and assessed from c_A.

On this semantics, we no longer get the result that $\ulcorner FAK^I_{now} : \Phi \urcorner$ and $\ulcorner Might : \Phi \urcorner$ are logically equivalent. To see that they could not be, it suffices to notice that the latter is assessment-sensitive while the former is not. However, a weaker kind of equivalence holds: they are *diagonally equivalent*.[22]

> Two sentences Φ and Ψ are *diagonally equivalent* iff for any possible context c, Φ is true as used at and assessed from c just in case Ψ is true as used at and assessed from c.

That is, a speaker considering $\ulcorner FAK^I_{now} : \Phi \urcorner$ and $\ulcorner Might : \Phi \urcorner$ from a particular context c should hold that an occurrence of either at c would have the same truth value. This vindicates the intuition that it is correct to say "It is possible that p" just when what one knows does not exclude p.

6.4 Monadic "true"

The relativist semantics makes use of two relativized truth predicates: (1) truth of a formula at a point of evaluation and (2) truth of a sentence at a context of use and context of assessment. These are theoretical notions that get their significance from the role they play in a larger theory of meaning.[23] But what about the *monadic* predicate "true" used by ordinary speakers—a predicate that applies to *propositions*, not to sentences? Can the relativist make sense of this? Yes—it's just another bit of vocabulary in the object language, and we can give semantics for it just as we can for other predicates.

First we need a proto-theory of propositions. We don't need to say in any detail what propositions are, or how they are individuated. We will assume only that there *is* such a thing as (for example) the proposition that Smith might be the murderer. Supposing there is such a proposition, what can we say about its truth? It is standardly assumed that propositions will have different truth values relative to different possible worlds. But the relativist

[22] The proof is straightforward.

[23] For details, see MacFarlane 2005b.

will also take the truth of propositions to be relative to an information state, which we have been modeling by a *set* of worlds, those not excluded by "what is known."[24]

Of course, the object-language truth predicate does not have argument places for these; it is monadic. So an account of its semantics must explain how these argument places are to be filled in. The answer is obvious: we just extract these values from our points of evaluation.

> The extension of "True" at a point of evaluation $\langle c_U, w, i, a\rangle$ is the set of propositions p such that p is true at $\langle w, i\rangle$.

"True" so defined is disquotational: every instance of the following schema is logically necessary (true at every point of evaluation):

$$\forall x((x = \text{the proposition that } P) \supset (True(x) \equiv P))$$

(where P is replaced by a sentence).[25] This is a welcome result. A disquota-

[24]Note that this view about propositional truth is not what makes the relativist view "relativist," since it is compatible with a "nonindexical" form of contextualism (see MacFarlane forthcoming). Indeed, our Solipsistic Contextualist could embrace the idea that truth for epistemic modal propositions is relative to a world and an information state. The crucial issue between the contextualist and the relativist is whether truth varies with the context of assessment, and that is left open by this decision about propositional truth.

Egan et al. 2005 and Egan 2007 take truth for epistemic modal propositions to be relative to a world, time, and an individual as "center." This approach may seem simpler than the one proposed here, and less radical, so it is worth taking a moment to explain why I am not inclined to go this way. The basic problem is that, although an world-time-individual triple is guaranteed to determine a pair of a privileged world and information state—the set of worlds not excluded by what is known by the individual at the world and time—the reverse is not the case. Given an arbitrary world w and information state i, there is no guarantee that there will be a triple $\langle w, t, s\rangle$ such that i is the set of worlds not excluded by what is known by s at w and t. Indeed, we know that some combinations of i and w will not be determined by *any* $\langle w, t, s\rangle$. For, knowledge being factive, w must surely belong to the set of worlds not excluded by what is known by s at w and t, so centered worlds will not determine any $\langle w, i\rangle$ pairs where $w \notin i$.

Why does this matter? Well, suppose that at c_A we are assessing an assertion at c_U of the proposition that p. We should judge the assertion true just in case p is true at $\langle w_{c_U}, i_{c_A}\rangle$. If we are assessing a merely counterfactual assertion, so that the world of c_U is not our world, it may be that $w_{c_U} \notin i_{c_A}$. This is no problem if we take propositional truth to be relative to world-set pairs. But what do we do if we take propositional truth to be relative to world-time-individual triples? There's not going to be a triple that gives us the world *and* the set of nonexcluded worlds we need. The problem, in short, is that centered worlds "entangle" parameters that need to be free to move independently in the semantic theory.

[25]Proof: Take any sentence P and consider any point of evaluation $\langle c, w, i, a\rangle$ such that 'x = the proposition that P' is true at $\langle c, w, i, a\rangle$. '$True(x)$' is true at $\langle c, w, i, a\rangle$ iff $a(x)$ is true at $\langle w, i\rangle$. But this is so iff P is true at $\langle c, w, i, a\rangle$, because $a(x)$ is the proposition expressed by P at c. So '$True(x) \equiv P$' is true at $\langle c, w, i, a\rangle$.

tional truth predicate is a useful expressive device, and it is reassuring that the relativist can make good sense of it.

It is a corollary of this result that when P is assessment-sensitive and α denotes the proposition expressed by P, $\ulcorner True(\alpha) \urcorner$ will also be assessment-sensitive.

7 Tensed Epistemic Modals

I want to close with a discussion of two problems I regard as open and difficult. The first concerns the interaction of epistemic modals and tense; the second concerns the robustness of the data used to motivate the relativist semantics. I will not try to resolve these issues here; the aim is to provide a prolegomenon to further investigations.

On all of the views we've considered so far, the set of epistemically open worlds with respect to which epistemic modals are evaluated is supplied entirely by context. For the Solipsistic Contextualist, it is the set of worlds not excluded by what the speaker knows at the time of utterance; for the Nonsolipsistic Contextualist, it is the set of worlds not excluded by what the contextually relevant group knows at the time of utterance; for the Solipsistic Relativist, it is the set of worlds not excluded by what the assessor knows at the time of assessment. Let us suppose all these contextual factors have been fixed. Then either the set of nonexcluded worlds contains worlds at which Fermat's Last Theorem is false or it does not. If it does, then the second conjunct of

(10) In 1980 it was possible that Fermat's Last Theorem was false, but this is not possible today.

is false. If it doesn't, then the first conjunct is false: the past tense has no effect, because there is no time variable associated with the epistemic modal. Either way, then, (10) is false. On all of these views, "possible" and "might" are temporally *rigid*: like "now" and "yesterday," they are unaffected by shifts in the time of evaluation.

This might be thought to be a bad consequence. I am not so sure. Note, first, that the information relevant to evaluating embedded occurrences of epistemic modals does not seem to shift with the *world* of evaluation. Here is a test case:

(11) It isn't possible that Jones is the murderer, but if no one had looked in this desk, it would have been possible that Jones was the murderer.

But if counterfactual changes in what we know do not induce counterfactual changes in what is epistemically possible, why should trans-temporal changes in what we know induce trans-temporal changes in what is epistemically possible? The (alethic) modal rigidity of epistemic modals is some evidence for their temporal rigidity. Intuitions about (10) are not as clear, but to me it has the same odd feel as (11).[26] A much more natural thing to say would be

(12) In 1980 people thought it possible that Fermat's Last Theorem was false, but we know today that this is not possible.

Here is a similar test case with a contingent sentence embedded under the epistemic modal:

(13) We know now that Sarah murdered Jenkins by herself. But yesterday it was possible that she had an accomplice.

We do sometimes hear this kind of thing, especially in legal contexts. But my (admittedly contaminated) intuition is that it is, strictly speaking, false. What is meant is

(14) We know now that Sarah must have murdered Jenkins by herself. But yesterday it was possible *for all we knew* that she had an accomplice. (Or: yesterday it had *seemed* possible that she had an accomplice.)

If you're not convinced, consider asking someone who asserts (13) *at what time* it became impossible that Sarah had an accomplice, and what changed to make it so. I predict embarrassment. For the only answer is, "At *N* o'clock, when we learned such and such." And this answer commits the speaker to the view that by learning something, she *made it impossible* that Sarah had an accomplice. I believe that ordinary speakers (those not already indoctrinated into contextualist theories of epistemic modals) will find this consequence bizarre.

Here's a real-life example, from a *New York Times* article concerning *Science*'s retraction of a paper reporting the production of eleven lines of cloned human embryonic stem cells:[27]

[26] If we cash in "possible" for "might," it sounds even worse: "In 1980 Fermat's Last Theorem might have been false, but today it must be true." But this may be due to syntactic differences between "might" and "possible."

[27] Gina Kolata, "Amid Confusion, Journal Retracts Korean's Stem Cell Paper," *The New York Times*, Saturday, December 31, 2005.

> The retraction did not include information revealed in South Korea at a news conference on Thursday. Until then, *it had seemed possible* that Dr. Hwang's group had created 2 cloned stem cell lines, not 11. On Thursday, the investigators in Seoul said that even those two were not clones. (emphasis added)

Notice how odd it would have been to say, "Until then, it had *been* possible that Dr. Hwang's group had created 2 cloned stem cell lines, not 11," despite the fact that the information revealing that no cloned lines had been created was not known until the press conference.

Thus it is far from clear that simple temporal embeddings like (10) and (13) should motivate us to make epistemic modals time-indexed. But there are more complex cases that are harder to dismiss in this way, for example:

(15) I studied that book because it was possible that Fermat's Last Theorem would be refuted using its techniques.

Intuitively, it seems that one can truly and felicitously assert (15) even if one knows that Fermat's Last Theorem is true, provided it was not known at the time one studied the book that the theorem was true. But "becausal" contexts are generally considered to be factive: "*A* because *B*" implies, or perhaps presupposes, *B*. So, if (15) is true and felicitous, then it seems we are committed to the truth of "it was possible that Fermat's Last Theorem would be refuted using its techniques," even though we know now that the theorem would not be refuted.

It would be hasty to conclude, however, that the set of worlds relative to which an occurrence of "possible" shifts with the time of evaluation. Consider this close analogue of (15):

(16) I failed to prove the conjecture because it was impossible to prove it.

Here the set of worlds over which "impossible" quantifies is definitely not the set of worlds left open by the speaker's knowledge at the time of evaluation (that is, the time the proof was being attempted).

It is presumably important here that (15), unlike (16), gives the speaker's own reasons for doing something. In this kind of context, the presupposition of factivity is sometimes relaxed. Suppose Joe has just found out that the internet search company he invested in has gone bankrupt. He might felicitously say:

(17) I bought that stock because it was going to be the next Google!

Similarly, the utterer of (15) might say:

(18) I studied that book because it was going to show me how to refute Fermat's Last Theorem.

Was the book going to show her how to refute Fermat's Last Theorem? Presumably not, since the theorem is true. But the speech is still felicitous, in a context where the speaker is displaying the mental state that motivated a certain action. Given examples like (17) and (18), which do not contain epistemic modals, (15) and sentences like it do not motivate taking epistemic modals to be temporally "shifty."

Another challenging class of cases involves *binding*:

(19) He lectures in a bulletproof vest whenever it is possible that members of the audience are packing handguns.

(20) Whenever it was possible that Mary was drunk, the people she came with drove her home.[28]

It seems that in (19), "possible" needs to be evaluated with respect to the sets of worlds left open by what is known (presumably by the lecturer) on various occasions of lecturing. And in (20), "possible" needs to be evaluated with respect to what is known on various occasions of partying.

Note that in addition to the bound readings of (19) and (20), there are readings where the occurrences of "possible" are, plausibly, assessment-sensitive. Suppose several groups of researchers have been compiling data on a lecturer's sartorial habits, and about the likelihood that members of various audiences are armed. The groups might disagree about whether, at some particular lecture where no bulletproof vest was worn, it was possible that the audience members were packing handguns, and because of this they might disagree about the truth of an occurrence of (19). They would *not* be disagreeing about whether the lecturer's knowledge left it open that the audience members were armed; they might all agree that it did. A similar scenario will generate an assessment-sensitive reading of (20). So it is not the case that the modals are *always* affected by tense (or, relatedly, quantification over events). Sometimes there is a kind of binding, and sometimes there is not. When there is not, the arguments for the assessment sensitivity of the modals apply as before.

Moreover, it is not just the time that gets bound, but the *knower*. This can be seen in the first of two possible bound readings of (20):

(21) Whenever it was possible [for all they knew then] that Mary was drunk, the people she came with drove her home.

[28] I owe this example to Fabrizio Cariani.

(22) Whenever it was possible [for all I knew then] that Mary was drunk, the people she came with drove her home.[29]

So the simple expedient of indexing the epistemic modals to a time (that is, letting the set of worlds they quantify over be a function of the time of evaluation) is not going to be sufficient to generate all the bound readings.

The most straightforward approach for the relativist—one that works within the framework that has been presented above—is to get the various bound readings by appeal to "free enrichment" of (19) and (20) with appropriately placed FAK_t^α : operators. These operators provide agent and time variables that can be bound by the enclosing quantifiers. The idea, then, is that one can use the *sentence* (20) to express a number of different propositions, including at least two assessment-invariant propositions with binding of the variables in FAK_t^α : operators, as well as one assessment-sensitive proposition with no binding. The speaker relies on the audience's ability to figure out which proposition she intends to convey based on contextual clues. (When the ambiguity cannot be resolved contextually, the FAK_t^α : operators can be made explicit.)

To sum up: The simpleminded approach of taking the modals to be sensitive to what is known at the time of evaluation generates dubious predictions, and is neither necessary nor sufficient to explain what is going on in sentences (19) and (20). It is not sufficient, because the binding involves not just times but knowers. And it is not necessary, because the binding can be explained by positing an implicit FAK_t^α : operator. Since this explanation is consistent with the relativist semantics, interactions between temporal modifiers and epistemic modals do not give us any compelling reason to resist taking epistemic modals to be assessment-sensitive.

8 Doubts about the Data

The second set of worries I want to discuss concerns the robustness of the data used to motivate the relativist semantics.

8.1 Limits to retrospective correction?

Immediately after presenting his "salvage ship" case, discussed in section 3.2, above, Ian Hacking writes:

[29] If you're having trouble getting this reading, try continuing (20) with "This seems to have been a lucky coincidence, because they were usually too plastered to notice her condition."

> When one starts collecting examples like this, it begins to look as if, whenever it turns out to be false that p, we say, of an earlier era, that in those times it may have seemed possible that p, but it was not really possible at all.

If Hacking had endorsed this description of the data, he would have been well on the road to relativism. For only a relativist semantics can explain why earlier epistemic modal claims are *always* evaluated in light of what we know now (at the time of assessment), even when we know much more than was known at the time the claim was made.

However, Hacking thinks that this description of the data "would be too strong." Here's why:

> Consider a person who buys a lottery ticket. At the time he buys his ticket we shall say it is possible he will win, though probably he will not. As expected, he loses. But retrospectively it would be absurd to report that it only *seemed* possible that the man would win. It was perfectly possible that he would win. To see this clearly, consider a slightly different case, in which the lottery is not above board; it is rigged so that only the proprietors can win. Thus, however it may have seemed to the gullible customer, it really was not possible that he would win. It only seemed so. "Seemed possible" and "was possible" both have work cut out for them. (Hacking 1967, 148)

If Hacking is interpreting his example correctly, it spells trouble for the relativist. For it suggests that the retrospective assessment data used to motivate relativism do not extend as far as the relativist needs them to. In the case of the non-rigged lottery, it seems, we *don't* assess our earlier claim that it was possible that the man would win as false, despite the fact that what we know now (after the lottery) excludes his having won. This seems to favor some version of contextualism over relativism.

However, it is far from clear that Hacking's interpretation of the example is correct. Hacking says,

(23) It was perfectly possible that he would win,

and this seems right. But assent to (23) is only problematic for the relativist if "possible" in it is an epistemic modal. And there are at least three reasons for supposing that it is not:

1. The embedded clause ("that he would win") is in the subjunctive mood. Epistemic uses of "possible" characteristically take the indicative. So,

let's try forcing an epistemic reading by putting the clause in the indicative (rephrasing it a bit to avoid grammatical difficulties): "It was perfectly possible that he had the winning ticket." Now my willingness to accept the sentence vanishes. We know he did not, in fact, have the winning ticket, so we can't assert that it was possible that he did.

2. Suppose the universe evolves deterministically. Does that assumption make a difference to your willingness to accept (23)? If it does—and it does for me—that is strong evidence that the modal in (23) is alethic. For whether the universe evolves deterministically is independent of the truth of epistemic modal claims. Determinism is compatible with universal ignorance about how things will evolve.

3. We will certainly not accept "It *is* perfectly possible that his ticket was going to be the winning one." So if we accept (23), we will have a case like (13), which we argued was strictly speaking false.

If "possible" in (23) is not an epistemic modal, then Hacking's example does nothing to call into question the evidence supporting an assessment-sensitive semantics for epistemic modals.

8.2 Ignorant assessors

Dietz 2008 has observed that although our intuitions about retrospective assessments seem to support relativist semantics when the assessor knows *more* than the original asserter, they do not do so when the assessor knows *less*. Here is a variation on one of Dietz's examples. Suppose that yesterday I proved Theorem X and asserted "Theorem X must be true." Today, however, my memory has gone fuzzy. I recall that I was working on Theorem X, but I don't remember whether I proved it, refuted it, or did neither. If Solipsistic Relativism is correct, I should be able to say:

(24) If I said "Theorem X must be true" yesterday, then what I said was false.

For what I know now (at the context of assessment) leaves open the possibility that Theorem X is false.[30] And this seems bizarre. Intuitively, I don't have warrant to pronounce on the falsity of claims made by my better-informed past self, even when these claims contain epistemic modals.

[30]Note that given the account of "True" from section 6.4, above, it does not matter whether we say "is false" or "was false" in (24).

If epistemic possibility is perspectival, this data suggests, it is *asymmetrically* perspectival. The truth of epistemic modal claims can depend on what is known by the assessor, but only if the assessor knows more than the original asserter.

One way to capture this asymmetry in the relativist account would be to complicate the semantics, amalgamating the asserter's and the assessor's knowledge into a single body of known facts with respect to which the epistemic modal is to be evaluated. The definition of truth of an occurrence of a sentence in context would then look like this (only the second bulleted item has changed from section 6.3):

> An occurrence of a sentence Φ at a context c_U is true as assessed from a context c_A iff Φ is true at every point of evaluation $\langle c_U, w_{c_U}, i_{c_U+c_A}, a\rangle$, where
>
> - w_{c_U} = the world of c_U,
> - $i_{c_U+c_A}$ = the set of worlds that aren't excluded either by what is known (at c_A) by the agent of c_A or by what is known (at c_U) by the agent of c_U,[31]
> - a = an assignment of objects from the domain of c to the variables.

The revised account, which is a kind of HYBRID between Solipsistic Relativism and Solipsistic Contextualism, would agree with Solipsistic Relativism on every case where the assessor is not ignorant of any relevant facts that the utterer knows. This includes all of the cases we used to motivate the relativist account.

One consequence of the move to the hybrid account is that it makes it difficult to *reiterate* epistemic modal claims. Suppose Sally says (at time t)

(25) It's possible that Joe is six years old.

How can we make a claim with the same truth conditions as Sally's—one that is guaranteed to have the same truth value as hers relative to every context of assessment? For the Solipsistic Relativist, this task is easy:

(26) It's possible that Joe is six years old at t

[31] Another alternative would be to consider what is known *distributively* by the speaker at c_U and the assessor at c_A: $i_{c_U+c_A}$ = the set of worlds that aren't excluded by what would be known by a rational agent who knew everything known at c_A by the agent of c_A and everything known at c_U by the agent of c_U.

will do the trick. For the Solipsistic Contextualist, it is equally easy:

(27) For all Sally knows at t, Joe is six years old at t.

But on the hybrid account, neither of these sentences can be used to make a claim that can be counted on to have the same truth value as Sally's claim relative to every context of assessment. For the truth of (26), as uttered by us and assessed by Judy, will depend in part on what is known by us, while the truth of (25), as uttered by Sally and assesed by Judy, will not depend at all on what is known by us. And since (27) is not assessment-sensitive, it cannot serve to reiterate Sally's assertion of (25), which is. Thus, on the hybrid account, epistemic modal claims are perspectival in the very strong sense that a claim made from one perspective cannot be reiterated in another, except by the use of semantic ascent or anaphoric devices.

A more promising response to the objection—and one less concessive to contextualism—would be to hold that whose knowledge is relevant to the evaluation of epistemic modals is *itself* determined by features of the context of assessment.[32] We might call this more flexible and less commital view FLEXIBLE RELATIVISM:[33]

> An occurrence of a sentence Φ at a context c_U is true as assessed from a context c_A iff Φ is true at every point of evaluation $\langle c_U, w_{c_U}, i_{c_A}, a\rangle$, where
>
> - w_{c_U} = the world of c_U,
> - i_{c_A} = the set of worlds that aren't excluded by the information that is relevant at c_A,
> - a = an assignment of objects from the domain of c to the variables.

This is still a relativist view, because it is features of the context of assessment, not the context of use, that determine which information state is relevant for the evaluation of epistemic modals. But it refrains from making any additional commitments about what information is relevant. If it were always the information possessed by the agent of c_A, Solipsistic Relativism would be correct; if it were always a combination of the information possessed by the agents of c_A and c_U, the hybrid view would be correct. Flexible Relativism refrains from hard-wiring either choice into the semantics,

[32] The following section was added after Dietz 2008 was in press.

[33] Compare Flexible Contextualism, section 6.2, above.

holding instead that the choice will depend on features of the context of assessment.

This response is analogous to Keith DeRose's defense of epistemic contextualism[34] in the face of cases where epistemic standards seem to depend on the situation of the *subject* of the knowledge attribution, not the attributor. DeRose notes that the contextualist account, properly understood, can handle such cases without modification:

> ...there is nothing in contextualism to prevent a speaker's context from selecting epistemic standards appropriate to the subject's practical situation, even when the subject being discussed is no party to the speaker's conversation—which is good, because speakers often do select such standards when their conversational purposes call for it. On contextualism, the speaker's context does always call the shots. ...But sometimes speakers' own conversational purposes call for employing standards that are appropriate to the practical situation of the far-away subjects they are discussing, and so the shot that the speakers' context calls can be, and often quite naturally will be, to invoke the standards appropriate to the practical situation faced by the subject being discussed. (DeRose 2005, 189)

The essential point here is that it is the *speaker's* context that determines whether it is appropriate to take into account the situation of the subject. So the view is genuinely contextualist, not a hybrid of a contextualist and a subject-centered view.

Flexible Relativism affords a similar kind of response to worries about ignorant assessors. The idea is that, although in some cases the speaker's information *is* relevant to the evaluation of epistemic modal claims, it is the *assessor's* context that determines when this (or any other information) is relevant. In contexts where the primary point of the assessment is critical evaluation of a speaker's assertion (as when one is trying to determine whether the speaker might be a trustworthy source of information), the relevant information state will generally be a composite of the speaker's and the assessor's information. And in contexts where the assessor is simply trying to guide her own inquiry, the relevant information state may be entirely determined by her own knowledge. But in each case, it is features of the context of assessment that determine which information is relevant. So the view is genuinely relativist, and not a hybrid of relativism and contextualism.

[34]The view that the extension of "knows" depends in part on contextually determined epistemic standards.

9 Conclusion

None of the standard accounts of epistemic modals works very well. Contextualist accounts can't make sense of retrospective assessments, retractions, and disagreement, no matter how much contextual flexibility they introduce; expressivist accounts flounder on some of the same data and require a Ptolemaic account of embedded uses of modals. Each view can be motivated by pointing to the shortcomings of the other, but neither is very satisfying in its own right.

Relativism looks like a promising alternative to these standard views. It seems to explain all of the data that motivate the others views, but it can also handle the problem cases that they can't handle.

Substantial problems remain. There are many difficult issues that arise in the compositional semantics—for example, concerning the interaction of epistemic modals with temporal modifiers—but these face all of the standard accounts as well. The *special* problems for the relativist are philosophical problems—for example, the problem of making sense of the assessment-relative truth predicate. Those have been kept off the table here. The present paper is an advertisement for their importance, not just for philosophers, but for natural language semanticists of all stripes.

References

Simon Blackburn. *Spreading the Word*. Oxford University Press, Oxford, 1984.

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

Keith DeRose. The ordinary language basis for contextualism, and the new invariantism. *Philosophical Quarterly*, 55:172–198, 2005.

Richard Dietz. Epistemic modals and correct disagreement. In Manuel Garcia-Carpintero and Max Kölbel, editors, *Relative Truth*. Oxford University Press, Oxford, 2008.

Andy Egan. Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133:1–22, 2007.

Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In G. Preyer and P. Peter, editors, *Contextualism in Philosophy*. Oxford University Press, Oxford, 2005.

Ian Hacking. Possibility. *Philosophical Review*, 76:143–168, 1967.

R. M. Hare. Some alleged differences between imperatives and indicatives. *Mind*, 76:309–326, 1967.

John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, Oxford, 2004.

David Kaplan. Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Joseph Almog, John Perry, and Howard Wettstein, editors, *Themes from Kaplan*, pages 481–566. Oxford University Press, Oxford, 1989.

David Lewis. Index, context, and content. In Stig Kanger and Sven Öhman, editors, *Philosophy and Grammar*. Reidel, Dordrecht, 1980.

John MacFarlane. The assessment sensitivity of knowledge attributions. *Oxford Studies in Epistemology*, 1:197–233, 2005a.

John MacFarlane. Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105:321–339, 2005b.

John MacFarlane. Relativism and disagreement. *Philosophical Studies*, 132: 17–31, 2007.

John MacFarlane. Nonindexical contextualism. *Synthese*, forthcoming.

G. E. Moore. *Commonplace Book, 1919–53*. Allen and Unwin, London, 1962.

Huw Price. Does 'probably' modify sense? *Australasian Journal of Philosophy*, 61:396–408, 1983.

Huw Price. Semantic minimalism and the Frege point. In S. L. Tsohatzidis, editor, *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*, pages 132–155. Routledge, 1994.

Jason Stanley. Fallibilism and concessive knowledge attributions. *Analysis*, 2005.

Alfred Tarski. The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4:341–376, 1944.

Paul Teller. Epistemic possibility. *Philosophia*, 2:303–320, 1972.

Richmond H. Thomason. Indeterminist time and truth-value gaps. *Theoria*, 36:264–281, 1970.

Kai von Fintel and Sabine Iatridou. Epistemic containment. *Linguistic Inquiry*, 34:173–198, 2003.

Seth Yalcin. Epistemic modals. *Mind*, 116:983–1026, 2007.