

Relativism about epistemic modals

Andy Egan
Rutgers University
January 31, 2010

The paradigmatic instances of epistemic modals are the sorts of uses of “might”, “must” and “possible” that occur in exchanges like the following:

“Where are my keys?”

“I’m not sure – they might be on the desk.”

“Where’s Bob?”

“He must be in his office. His light’s on, and his jacket’s hanging outside the door.”

“Does John have cancer?”

“It’s possible that John has cancer. He has some of the symptoms. They’ve run some tests, but we won’t know the results until tomorrow.”

These are uses of “might”, “must” and “possible” whose truth is somehow or other bound up with something about somebody’s epistemic state. Call these kinds of sentences, in which an epistemic use of “might”, “must” or “possible” takes wide scope over an unmodalized clause, *simple epistemic modal sentences*, or *simple EM sentences*.

The overwhelmingly natural first thing to say about just how such sentences are bound up with our epistemic states – one that serves as the point of departure for most of the views about epistemic modals now in circulation in the literature – is that with these sorts of uses of “might”, “must” and “possible”, a simple epistemic modal sentence of the form *might:P* (or *possibly:P*) is true iff P is compatible with what’s known, and *must:P* is true iff not-P is incompatible with what’s known. The reason why epistemic modals are interesting, and why it’s hard to give a satisfactory theory of them, is that it’s remarkably difficult to transform that plausible-looking first shot into a worked out account.

The first obvious question to ask about such an account is, “known by whom”? And a brief survey of cases reveals that the answer has got to be, “it depends”. A number of authors have presented a variety of cases that demonstrate this. (See for example Hacking 1967, Teller 1972, Kratzer 1986, DeRose 1991) Here is a somewhat simplified version of one from DeRose (1991):

The Cancer Test

John’s just had a test, a negative result on which would rule out cancer. Jane, not knowing what the results were, says (in response to an inquiry from Bill, who’s heard a rumor that John has cancer), “It’s possible that John has cancer – they’ve run a test that could rule it out, but we won’t know the results until tomorrow”. At the same time, the doctor, who has seen the negative results of the test, says “it’s not possible that John has cancer – we should start planning tests for other diseases”.

It seems as if both Jane and the doctor speak truly. A plausible explanation of this is that it’s Jane’s knowledge that’s relevant to the truth or falsity of Jane’s claim, and it’s the doctor’s knowledge that’s relevant to the truth or falsity of the doctor’s claim. It’s easy to construct similar examples for “might”. (Though it’s complicated a bit by the inconvenient fact that there is no lexical item that clearly stands to “might” as “impossible” stands to “possible”, which makes the examples a bit less clean.)

The lesson that’s typically been drawn from these kinds of cases is a *contextualist* one. Different utterances of epistemic modal sentences are responsive to different peoples’ epistemic states because they are *about* different people’s epistemic states. At a first pass: If Jane says “John might have cancer”, she speaks truly, because what she’s said is *that it’s compatible with what Jane knows that John has cancer*. If the doctor says “John might have cancer”, he speaks falsely, since what he says is *that it’s compatible with what the doctor knows that John has cancer*.

This is actually not quite the standard contextualist conclusion. The

standard view has it that the relevant epistemic state needn't be the state of some *individual*, but could be (and often is) that of some group or community.

One way to motivate this is by drawing attention to a phenomenon observed by G.E. Moore (1962), in which he notes that it's possible to deny someone else's assertion of "it's possible that Hitler is dead by now" by saying, "I know he's not". It would be odd if an assertion about *my* knowledge could serve to deny your epistemic possibility claim if what epistemic possibility claims assert is just that the relevant proposition is compatible with *the speaker's* knowledge. It's to be expected, though, if what's being asserted is that it's compatible with the collective knowledge of a group to which we both belong. (There are complications about how to understand collective knowledge. See von Fintel and Gillies (forthcoming) for discussion.)

Another motivation comes from a family of cases discussed by Keith DeRose (1991), which are variations on *The Cancer Test* above. In one such case, Jane knows that the doctors have looked at the test results, but they haven't yet told her about them. Asked, "is it possible that John has cancer?", she replies, "I don't know whether it's possible that John has cancer; only the doctors know. I'll find that out tomorrow when the results of the test are revealed." Here, it seems clear that the embedded possibility claim can't be about her knowledge – what she doesn't know is whether it's compatible with the doctors' knowledge that John has cancer, so the relevant group must include at least the doctors.

There is another family of examples that motivates variability in whose knowledge is relevant, largely inspired by (or anyway in the spirit of) a footnote in John Hawthorne's *Knowledge and Lotteries* (Hawthorne 2006). These examples are meant to support the more radical hypothesis that a *single utterance* can be correctly evaluated as true by one assessor, and correctly evaluated as false by another, because the epistemic states that are relevant to correct evaluation are different for the two assessors. These kinds of examples figure prominently in the relativist arguments of, for example, MacFarlane (forthcoming), Egan, Hawthorne and Weatherson (2005), and Egan (2007).

Here is the relevant footnote:

[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 2006, p29, fn69)

And here is a case of the sort used in the relativist literature. (This one is a modified version of one from MacFarlane (forthcoming)):

Eavesdropping Brian

The evidence Sally and George have available to them leaves it open that Joe is either in Boston or in Berkeley. Brian, listening in from behind a nearby shrub, just saw Joe in Berkeley five minutes ago. Sally says to George, "Joe might be in Boston". George accepts Sally's assertion, and signals his assent by saying "that's true." Brian rejects Sally's assertion, and signals his dissent by muttering to himself, "that's false".

George's assent, and his attribution of truth to Sally's utterance, seem completely in order. So do Brian's dissent, and his attribution of falsity to the very same utterance. This suggests that it's not just that different epistemic states can be relevant to the evaluation of different utterances of epistemic modal sentences, but that different epistemic states can be relevant to different assessors' evaluations of the *same* utterance of an epistemic modal sentence. (As mentioned above, these cases are much more controversial than the cases meant to show variation in

relevant epistemic state across utterances. We'll talk more about the objections and responses to these kinds of cases, and the arguments based on them, in section five.)

Pretty much all parties to the debate agree that we can't give a uniform, one-size fits all answer to the question of whose knowledge is relevant to assessing the truth or falsity of epistemic modal claims. Many also question whether it's really always *knowledge* that's at issue, or whether we should allow the relevant epistemic relation to vary as well. Here is a case from Hacking (1967)

The Salvage Ship

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 some mistakes 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 thirty 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. (Hacking, 1967, p148)

Cases such as Hacking's suggest that we sometimes want a relation *weaker* than knowledge, so that sometimes what's relevant to the truth or falsity of an epistemic modal claim is not what we *know*, but, for example, what we could come to know by some available method of inquiry. (Other cases – some of which are discussed in Dever (forthcoming) – suggest that we sometimes want a relation *stronger* than knowledge, such that not everything that we know will always be relevant.)

In general, the truth or falsity of epistemic modal claims depends on the compatibility of the embedded proposition (the *prejacent*) with some body of information – the information that's within the epistemic reach of some person or group. Which person or group's epistemic reach is relevant, and what it takes for

some piece of information to count as within a person's or group's epistemic reach, seems to be variable, varying at least with the context in which epistemic modal claims are made, and possibly also with the context in which they're assessed. This is why it's so difficult to turn the promising first shot at the truth-conditions of epistemic modal claims – that e.g. *might:P* is true iff P is compatible with what's known – into a worked out theory.

One way to deal with this, and to allow for variation in which body of information is relevant, is to go contextualist, and to say that the relevant body of information is determined by the situation of the *speaker* – by the person who's uttering the sentence. Another is to go *relativist*, and say that the relevant body of information is determined by the situation of the *assessor* – by the person who's assessing the sentence for truth or falsity.

Our main focus here will be on relativism, and my goal here will be to put on clear display the outlines of the debate about relativism about epistemic modals. But it will be helpful to say a bit more about the structure of contextualist theories, since contextualism is the main competitor to relativism, and probably is (and ought to be) the default starting-point view. Accordingly, much of the motivation for relativism comes from the purported inadequacy of the contextualist options.

Two notes about things I won't do in what follows, before we move on:

First, there are (at least) two other options, which I won't discuss here for reasons of space: One option that's very much alive is to go with a theory according to which epistemic modal claims aren't in the truth and falsity business. (For advocacy of such proposals, see Yalcin 2007 and forthcoming, Swanson forthcoming.) Another thing to notice is that both relativist theories and standard contextualist theories assume quite a close connection between semantic content and communicative import. We could go with a theory that does separates these quite sharply, as do the sorts of *pluralist* theory offered (though not specifically about epistemic modals) by Cappelen and Lepore (2005), or the sort of propositional radical theory offered by Bach (forthcoming).

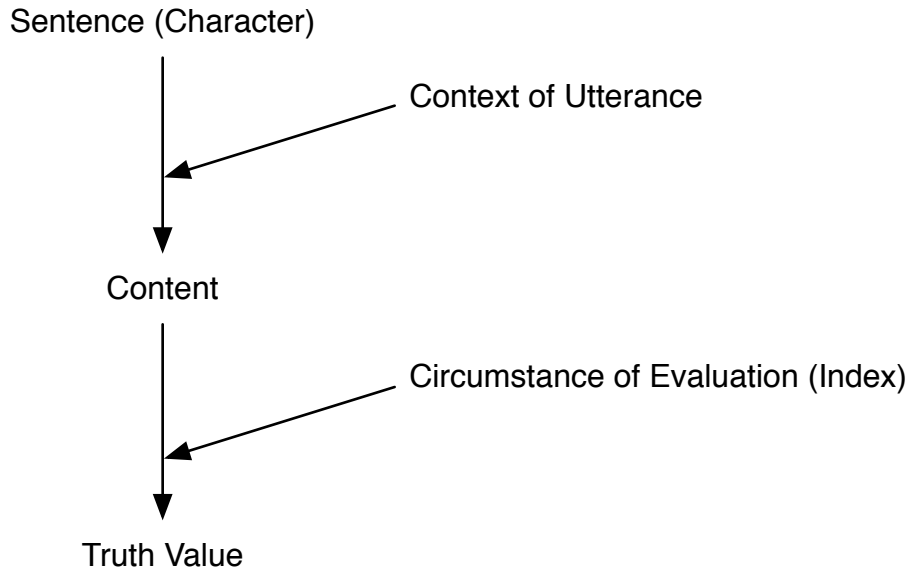
Second, I'm going to follow the literature by focusing almost entirely on “might” and “must” - and much more heavily on “might” than “must”. But it's worth

bearing in mind Eric Swanson's (forthcoming) warning that "might" and "must" are just two examples of a much broader category – the category Swanson calls "the language of subjective uncertainty" – which it would be nice to have a unified account of, and that there are theoretical risks associated with focusing too hard on a limited range of examples. I will run these risks here in the interests of space and ease of presentation, but it's worth bearing them in mind as we move along.

In the next section, we'll look at some of the important features of contextualist views in general. We'll then look, in section three, at contextualist views about epistemic modals in particular. In section four, we'll discuss the internal workings of two different sorts of relativist theories. In the fifth section, we'll discuss some standard arguments relativists have deployed to motivate relativism over contextualism, and we'll trace a few of the first steps of the ensuing dialectic of reply and objection. In that section, we'll look at a number of the standard replies to relativist arguments. At more or less every point in what follows, I will, unavoidably, leave out a lot. Rather than trying to lay out all of the fine details of the various branches of the dialectic, I'll be concerned primarily to lay out the central moving parts of the main positions in debate and to walk through the first few steps of the core arguments.

2. Contextualism

The following diagram illustrates the more-or-less standard picture of how sentences and utterances come by their truth-values. (The canonical source of this picture is, of course, Kaplan 1989. See also Lewis 1980.)



A sentence is associated, in the first instance, with a *character*, which determines a function from contexts of utterance to contents expressed. (It's more standard to talk about the characters of lexical items than of sentences. But it's easy to assign a character to a sentence based on the characters of its constituent expressions.) A sentence that doesn't include any context-sensitive vocabulary will have a constant character – it will express the same content in every context of use.

Assigning truth-values to sentences in a context of utterance is a two-step process: first, the context of utterance serves as an input to the sentence's *character*, in order to determine a content. Then the *circumstance of evaluation (or index)* corresponding to the context of utterance serves as the input to the contextually determined content, in order to determine a truth-value. So the truth-value of a sentence in a context, or of a particular token utterance, is sensitive to the context of utterance twice over: first, as an input to character, and second, as the determinant of which index serves as the relevant input to content.

The key feature of all of this for our purposes is that standard contextualist views deliver a single, once-and-for-all verdict on the truth-value of a sentence in a particular context of use, and (therefore) on the truth-values of particular token

utterances. While different utterances of the same sentence can take different truth-values, a particular, dated, world-bound utterance gets a truth-value *simpliciter*.

(I am going to assume that it's safe to – as I did above – move smoothly between issues about the assignment of truth-values to *sentences in context*, and issues about the assignment of truth-values to *particular token utterances*. I should note that this is potentially contentious.)

3. Contextualism about Epistemic Modals

There is a widely accepted, received view about the semantics of modals in general, the paradigmatic statement of which is found in Kratzer (1977 and 1981). Here is, I think, a pretty uncontroversial, minimally formal way of stating it: The truth of *modal:P* at a world *w* depends on how things stand with respect to *P* in the worlds *w** that bear some particular relation to *w* – the *w**s *accessible* from *w*. (More generally, the truth of *modal:P* at an index *i* depends on how things are with respect to *P* at the indices accessible from *i*.) Different species of modality are distinguished by differences in the relevant accessibility relation. So what's nomologically possible in *w* is what's true in some world that's *nomologically accessible* from *w* – some world whose history is compatible with the laws of *w*. What's nomologically necessary in *w* is what's true in every world whose history is compatible with the laws of *w*. What's epistemically possible in *w* is what's true in some world that is, in some sense, *epistemically accessible* from *w* – plausibly, some world that's compatible with the epistemic state of some selected person or group in *w*. And so on.

(The above is not quite right: for a number of modals, we need not just an accessibility relation, but something that imposes an ordering on worlds. (See Kratzer 1981, 1991.) We will ignore this complication in what follows.)

In this sort of framework, the distinction between metaphysical, nomological, deontic, etc. modals isn't a distinction between different sorts of lexical items, but between different kinds of uses of the same lexical items. There's just a single "might" in the lexicon, but it's context-dependent, since the relevant accessibility relation is different in different contexts. Similarly for "must", "possibly", etc.

The overwhelmingly natural thing to say about *epistemic* modals, given this framework, and the thing we should say unless we're forced out of it, is that they are modals like any other, and context-dependent in just exactly the way that modals in general are context-dependent. Epistemic modals, that is, are just standard garden-variety modals, which happen to occur in a context that provides a distinctive sort of accessibility relation – one that has got something to do with somebody's epistemic state, such that the accessible worlds are the ones that some person or group stands in some particular epistemic relation to.

The simplest sort of contextualist account is probably what MacFarlane (forthcoming) calls *solipsistic contextualism*, according to which the accessible worlds are those compatible with the speaker's knowledge. It's worth noting that this is a theory that makes decisions at two different choice points. It says that it's the *speaker*, rather than some other person or group, whose epistemic state is relevant. That's captured by calling it *solipsistic* contextualism. It also says that the epistemic relation that's important is *knowledge*, rather than some other – possibly weaker, possibly stronger, possibly cross-cutting – epistemic relation. Different solipsistic contextualisms could differ on this point. So the view in question is probably better described as *solipsistic knowledge contextualism*.

The cases of variability we looked at in section 1 suggest that solipsistic knowledge contextualism can't be right. DeRose's "might be possible" cases, and the facts about disagreement, suggest that the *solipsistic* part won't do, and cases like Hacking's suggest that the *knowledge* part won't do, either.

It should be clear, though, that solipsistic knowledge contextualism isn't the only option available for a contextualist about epistemic modals. There's a *lot* of room for variation in the story about just which people or groups, and which epistemic relations, are relevant, and in which circumstances. Pick any group, and any epistemic relation, that you like, and there's an accessibility relation that tracks what's compatible with what that group stands in that epistemic relation to. There is a lot to say about contextualism-internal debates about just which persons, groups, and epistemic relations are potentially relevant, and just how the relevance of a person, group or epistemic relation depends on other features of the context.

But what's important for our purposes is just the fact that contextualism offers quite a versatile framework for fleshing out the details of a theory of epistemic modals.

On our working contextualist theory, there will be two relevant contextually variable moving parts: The relevant group, and the relevant epistemic relation. An utterance of "Bob might be in his office" will always express a proposition of the type, *that it's compatible with all of the evidence that G stands in R to that Bob is in his office*, but it will express different such propositions in different contexts of utterance, because different Gs and Rs will be relevant in different contexts. The proposition that's expressed, that is, is a proposition about some person or group's epistemic state. It's true in worlds where the contextually relevant group fails to stand in the contextually relevant epistemic relation to evidence that rules out Bob's being in his office.

To sum up: this is a view according to which *epistemic* modals are just regular, garden-variety modals, used in a context that provides a certain distinctive kind of accessibility relation. What's distinctive about them is that their accessibility relations are ones that track what some person or group stands in some epistemic relation to. This allows for a lot of internal diversity within the category of epistemic modals, depending on whose epistemic reach the accessibility relation is tracking, and on what kinds of standards of *reach* are being applied.

3. Relativist Proposals

There are a number of different relativist views of epistemic modals on offer. All of them are committed to denying that utterances of simple epistemic modal sentences get truth-values "once and for all", and maintaining that a single, dated utterance can be true relative to one context of assessment and false relative to another.

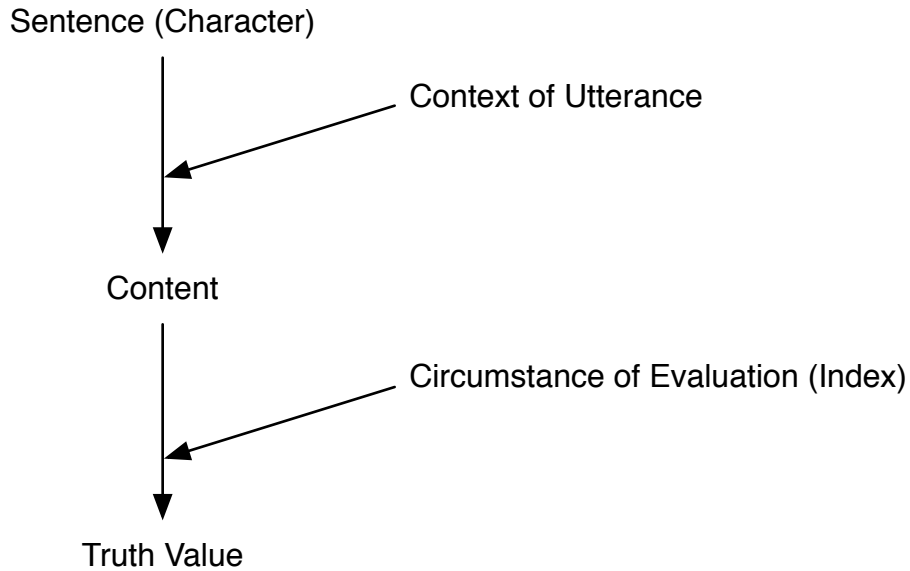
There are two ways to do this. First, one can say that the context of assessment plays a role in the determination of the *content* of the utterance, so that a single utterance can express different propositions to different audience members. Call such views *content-relativist*. (See Predelli 1996, 1998a, 1998b, Cappelen 2008a, 2008b, and Egan 2009 for defenses of such views – though not about

epistemic modals.) Alternatively, one can say that the utterance has a single determinate content, fixed by its context of utterance, but that (a) the content determines a truth-value relative to an index that includes not just a world, but also some further parameter with respect to which two contexts of assessment (within a possible world) might differ, and (b) which such index is relevant to assigning truth-values to an utterance can vary across contexts of assessment. Call such views *truth-relativist*.

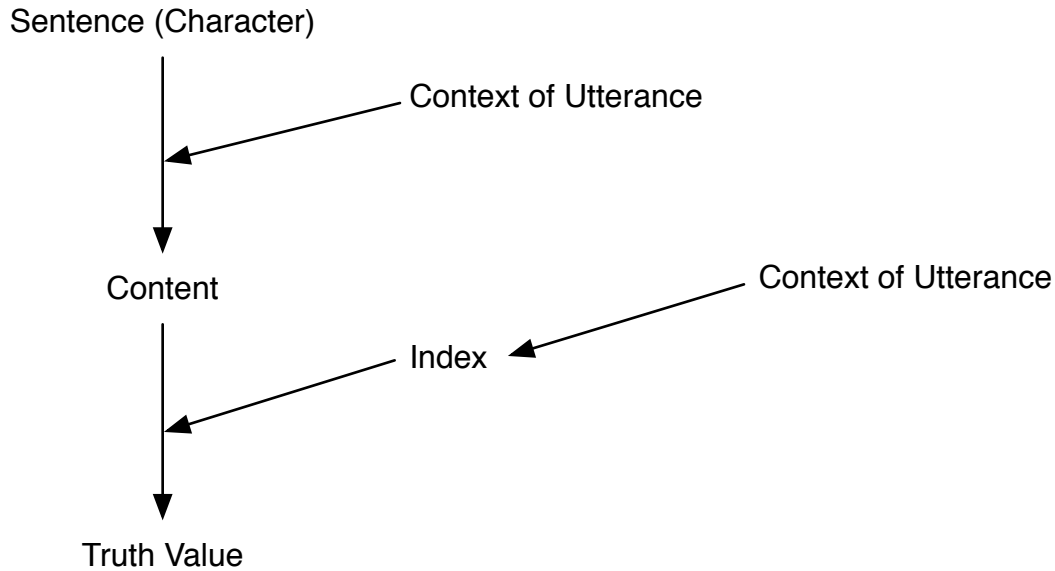
I'm going to set content-relativist theories aside at this point, since the proposals about epistemic modals that have been discussed under the heading of "relativism" in the literature are predominantly truth-relativist proposals. So in what follows I'll use "relativist" to mean *truth-relativist*, and the relativist theories I'll be discussing here will all be instances of that type.

The various truth-relativist theories about epistemic modals all do two things, corresponding to (a) and (b) above: First, they postulate a further parameter, in addition to a world (and perhaps a time) in the index, or circumstance of evaluation, relative to which the contents of epistemic modal sentences take their truth-values. And second, they tell a story about the role this extra parameter plays in the notion of utterance truth, and in the practice of evaluating utterances for truth and falsity, such that we get the possibility of one assessor correctly evaluating an utterance as true, while another correctly assesses it is false.

John MacFarlane's relativism (MacFarlane forthcoming) takes the propositions expressed by epistemic modal sentences to be functions from (at least) <world, information state> pairs to truth-values. (I'll write these '<w,q>' – using 'q' as a variable for information states and leaving 'i' to be used later on when we've got *individuals* appearing in our indices.) So the extra parameter is an information state – represented as a set of worlds. He then goes on to say that the process by which utterance truth is determined is more complicated than on the standard Kaplanian picture. Recall how we portrayed the Kaplanian picture before:



Really, this diagram leaves part of the story out. On the standard Kaplanian picture, not all indices are equal with respect to their role in the determination of truth-values for utterances, or for sentences in context. *Utterance* truth, and the accompanying notion of the truth of a sentence *S* in a context *C* – is sensitive only to the verdict that the content of *S* in *C* delivers about a *particular* index – the one that corresponds to the context *C*. (Depending on what parameters we’ve got in our indices, this is the index whose world is the world of *C*, whose speaker is the speaker of *C*, whose standards of precision are the standards in effect in *C*, etc.) If we want our diagram to reflect the full Kaplanian story about the determination of truth-values of sentences in context, we should write:

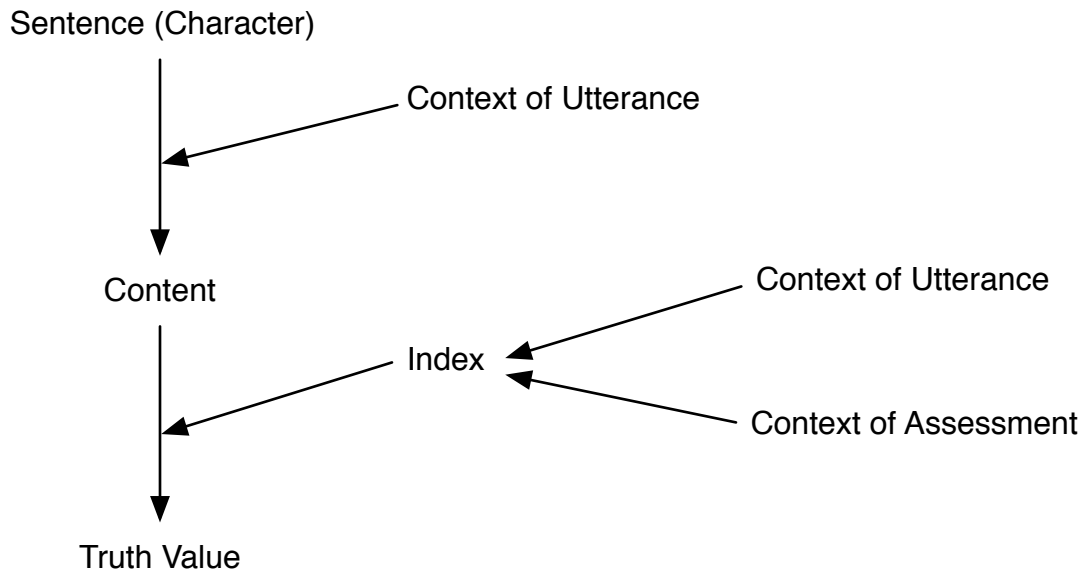


This addition allows our diagram to capture the way that, on the Kaplanian picture, the truth-value of a sentence in context is sensitive to context of utterance twice over – once as an input to character, and once as what we might call an *index-selector*. In particular, it indicates the special role that the context of utterance plays in the selection of the index that’s relevant to the *sentence’s* truth-value in the context of utterance, and so to the *utterance’s* truth-value *simpliciter*.

Making this explicit allows us to clearly display the difference between a MacFarlanian relativist account and a standard Kaplanian account. On MacFarlane’s sort of relativism, it’s crucial that, when assessing an utterance of S in C for truth or falsity, the determination of which index gets fed in to the content of S in C in order to determine the truth-value of the particular utterance is sensitive, not just to the context of utterance, but also to the context of assessment. (If index-selection was only sensitive to the context of utterance, then we’d be unable to get any variation of utterance truth across assessors, no matter how complicated we made our indices. See MacFarlane 2005, 2009.)

On MacFarlane’s picture, which index is relevant to the determination of utterance truth is determined partly by the context of utterance, and partly by the context of assessment. For example, ignoring times and assuming we have just

$\langle w, q \rangle$ indices, the relevant index for determining the truth of an utterance relative to an assessor is $\langle w_{CU}, q_{CA} \rangle$. (Where 'CU' names the context of utterance, and 'CA' the context of assessment.)



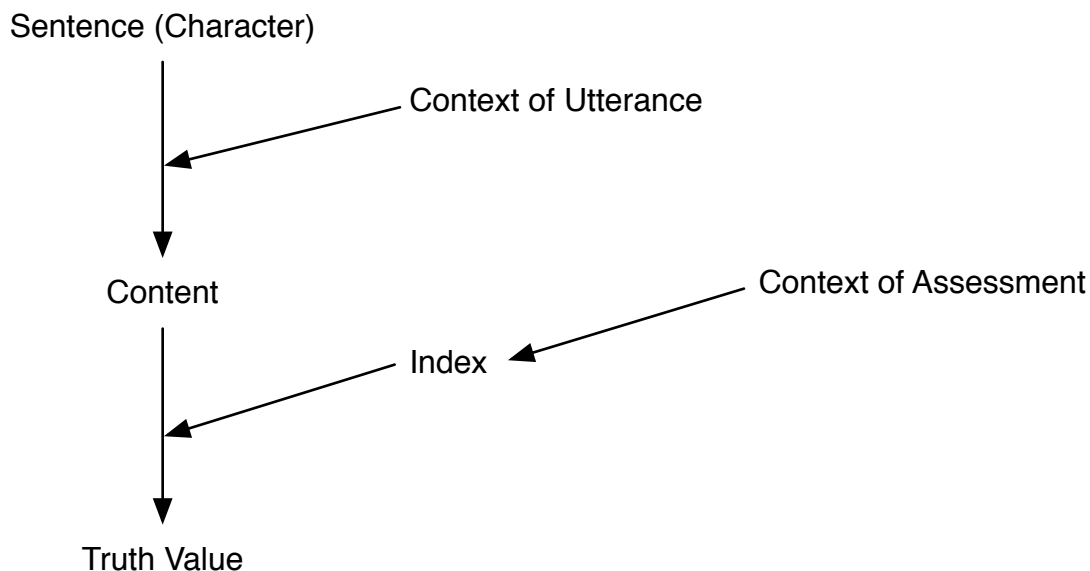
The important difference here, obviously, is that it's not the context of utterance alone that's responsible for selecting the index that's relevant to the assignment of truth-values to sentences in context, or to particular utterances. Index-selection, on this picture, is accomplished by the context of utterance and context of assessment together. As a result, this is a picture on which there's no such thing as the truth-value of a sentence in a context of utterance *simpliciter*, or of the truth-value of an utterance *simpliciter*. There's only the truth-value of a sentence relative to a *pair* of a context of utterance and a context of assessment, or of an utterance relative to a context of assessment. That's because it's only once we've got both a context of utterance and a context of assessment that we know which index to feed in to the content in order to get a truth-value. (On at least the first-pass version of MacFarlane's relativism about epistemic modals, *might:P* is true at $\langle CU, CA \rangle$ iff what is known to the assessor at CA is compatible with the truth of P at $\langle CU, CA \rangle$.)

This allows for a different kind of variation in whose informational state is relevant than we find in contextualist theories. Here, the truth-values of epistemic modal claims are sensitive to the information available to the *assessor*. As a result, we can get the same utterance correctly assessed as true by one assessor, and correctly assessed as false by another.

Another relativist option is a theory that trades in *de se* linguistic content, as advocated by Tamina Stephenson (2007) and myself (Egan, Hawthorne and Weatherson 2005, Egan 2007, forthcoming). Our views are different in detail, but for our purposes here I'll lump them together under the single heading of "*de se* relativism". On this kind of view, the parameter-adding is achieved by moving to a framework in which sentences in context have *de se* contents. (For more on the *de se*, see for example Lewis 1979, Chisholm 1981, Chierchia 1989, Feit 2008.) This is a picture of content on which the contents of sentences in context divide a space of possible *situations* or *predicaments*, rather than a space of worlds. Possible predicaments are standardly modeled with centered worlds, which we can think of as <world, time, individual> triples. (Peter Laserson's (2005, 2009) version of relativism is also in many ways similar, as is Max Kölbel's (2002, 2003, 2009), though their accounts are targeting other subject matters.)

Possible predicaments (or centered worlds) differ from one another in many ways, including differing with respect to the evidential situation of the individual at the center. Epistemic modal sentences express centered-worlds propositions that are true of all and only those in a certain kind of epistemic situation. (For example, *might:P* is true at all and only those <world, time, individual> triples <w,t,i> such that the evidence within i's epistemic reach at t in w doesn't rule out P.)

Once we're given this account of the nature of the indices, we have the first component of a relativist story: we have contents that can take different truth-values relative to your predicament and to mine. We still need the second part of the story: an account of the role that the context of assessment plays in index-selection, such that it can happen that different indices are relevant to the assignment of a truth-value of an utterance when assessed by different individuals at different times. The simplest picture is this:



This is again a picture according to which the fundamental notions are not truth-in-a-context-of-utterance *simpliciter* for sentences, or of truth *simpliciter* for utterances. For sentences, the fundamental notion is truth relative to a <context of utterance, context of assessment> pair, where the context of utterance acts as input to character, and the context of assessment plays the index-selection role. For utterances the fundamental notion is truth relative to a context of assessment.

This again makes the epistemic state of the assessor relevant to the truth and falsity of epistemic modal claims, in a way that it's not on the contextualist picture. And so it allows for variation in the truth-value of a single utterance relative to different assessors.

The above picture – on which the context of assessment always has sole responsibility for index-selection – is a natural one. But it's worth noting that it's not the only one available – this will be relevant later, when we're discussing some of the arguments for and against relativist theories.

One other thing that's worth noting briefly about the *de se* relativist view is that it can be stated in such a way that it's just an instance of the usual Kratzerian framework. The relativist needs to say that the relevant accessibility relations hold

between centered worlds, rather than worlds, and they need to draw a distinction between “coarse grained” accessibility relations such that accessibility depends only on the *world* members of the <world, time, individual> triples (and so would be equally well modeled by accessibility relations between worlds), and those “finer grained” accessibility relations that are sensitive to other members of the triple. The proposal will then be that one of the distinctive features of *epistemic* uses of modal expressions is that the relevant accessibility relation is fine-grained, while other types of modality are governed by coarse-grained accessibility relations.

Notice that the kind of view just discussed is one that is both contextualist and relativist. The relevant accessibility relation is still fixed by context. But there’s less diversity among epistemic accessibility relations contributed (at least in simple, unembedded environments) than standard contextualist views maintain.

4. Relativists’ Arguments Against Contextualism

Contextualist theories say that different utterances of the same sentence can take different truth values, because they express different contents. Relativist theories say that there’s more variation than this – that we can’t accommodate all of the truth-value variation that we see in terms of variation, across contexts of use, in which contents are expressed.

There are two major families of arguments for this: First, what I will call *common-content* arguments, which aim to show that two utterances which we’d like to say differ in their truth-value must, after all, have the same content (and so that variation in truth-value can’t be chalked up to variation in content). Second, what I’ll call *single-utterance* arguments, which aim to show that particular, token utterances can take different truth-values relative to different assessors. We’ll look at the common-content arguments first.

One family of argument for common content across occasions of utterance is based on *disquotational reporting*: If Fred says “Bob might be in his office”, it seems as if it’s systematically okay for Fran to deliver disquotational indirect speech reports and (assuming Fred was being sincere) disquotational belief reports such as, “Fred said that Bob might be in his office” and “Fred believes that Bob might be in

his office". This seems to be so regardless of what Fran's context is like – in particular, it's so regardless of whether Fred's and Fran's contexts are alike with respect to the features that will, given a contextualist account, plausibly be relevant to fixing the semantic value of "might". This isn't what you'd expect, at least at a first pass, given a contextualist account of "might".

Contextualist accounts seem to predict that "might" will take its semantic cues in the original utterance from *Fred's* context, and in the report from the *reporter's* context. If that were so, disquotational reports would run the same kind of risk of misreporting as do disquotational reports involving such paradigmatic indexicals as "I" and "here". If Fred says "I am in New Brunswick", or "it's raining here", it *isn't* systematically okay for Fran to deliver disquotational reports. Reporting with "Fred said that I am in New Brunswick" is systematically bad, and "Fred believes that it's raining here" is only okay if Fred's and Fran's contexts of utterance are alike with respect to the semantic value they determine for "here".

In general, the problem is that Fred's utterance in C1 expresses the proposition *that Bob might_{C1} be in his office*, while Fran's report in C2 expresses the proposition *that Fred said that Bob might_{C2} be in his office*. Since $C1 \neq C2$, it's very likely that "might" will express something different in the two contexts, and so it's very likely that Fran's report will be mistaken. (One natural possibility: Fran will misreport Fred as having said that it's compatible with what *Fran* knows that Bob is in his office.)

But disquotational reports are, in fact, completely safe. So, the relativist argument goes, there must not be a danger of "Bob might be in his office" having different semantic values in the context of utterance and the context of reporting. And so "Bob might be in his office" must have the same semantic value across the board. And so contextualism about "might" must be mistaken. And if contextualism about "might" isn't correct, then the best way to explain the variation that we see in the truth-values of different "might" claims is by saying that the stable content of "might" claims can take different truth-values relative to different contexts of

assessment. (Versions of this argument occur in, for example, Kölbel 2002, and Egan, Hawthorne and Weatherson 2005.)

Another kind of disquotation-based argument for common content is based on disquotational attributions of agreement and disagreement. If Larry says “Bob might be in his office”, and Lisa says “It’s not the case that Bob might be in his office”, Liz can report on this by saying “Larry and Lisa disagree about whether Bob might be in his office”. Similarly if Larry and Lisa both say “Bob might be in his office”, Liz can report on this by saying, “Larry and Lisa agree that Bob might be in his office”.

These kinds of reports seem safe – their appropriateness doesn’t seem to be hostage to the facts about what Larry’s, Lisa’s, or Liz’s contexts are like. But we wouldn’t – at least at first glance – *expect* these kinds of reports to be safe if contextualism about “might” were true. For example, given a contextualist view of “might”, there ought to be a danger that, since Larry was speaking in C1 and Lisa in C2, Larry said *that Bob might_{C1} be in his office*, and Lisa said *that it’s not the case that Bob might_{C2} be in his office*. And these two claims, since they’re (quite likely) about the epistemic situations of different persons or groups, needn’t be in any conflict with each other. Worse, since “might” in Liz’s report will get its semantic value from *Liz’s* context, she’s in danger of reporting Lisa and Larry as disagreeing about something that, quite likely, neither of them has any views about at all – for example, about whether it’s compatible with what Liz knows that Bob is in his office.

Notice that the argument here isn’t that, on a contextualist view, disquotational disagreement- or agreement- reports are *always* defective. The argument is that, in fact, the relevant kinds of disquotational reports are *always* okay, and contextualism predicts that they’ll only be okay in certain narrowly constrained circumstances, in which the two parties’ contexts are aligned with respect to the semantic value they determine for “might”. Liz doesn’t, in fact, have to verify that Lisa and Larry are in relevantly similar kinds of contexts before she can, with confidence, report them as disagreeing/agreeing. It’s enough that she knows that they sincerely produced the relevant sentences. (Mark Richard (2004, 2009) makes these kinds of arguments in favor of a type of relativism in other domains, though he does not endorse relativism about epistemic modals.)

The arguments we've just canvassed are all arguments that the content of epistemic modal claims is stable across contexts of utterance, and so to the extent that we've got different intuitions about the truth-values of different utterances of a *might:P* claim, that must be because it's got the sort of content that can vary in truth-value at different points of evaluation within a world, rather than because it's expressing different contents in different contexts of use. Since the variation in truth-value across utterances isn't attributable to variation in *content*, it must be that the *same* content is taking different truth-values at the different indices relevant to assessment of the utterances.

This leads to a worry about all of these arguments: These look like better arguments for the sort of view MacFarlane (2009) describes as "nonindexical contextualism" than for any of the views that have been advocated under the heading of *relativism* about epistemic modals. They're arguments that different utterances of the same epistemic modal sentence (within a world) can take different truth-values, without differing in content. But they're not arguments that a *single utterance* of an epistemic modal sentence can take different truth-values relative to different assessors. And so they're not (at least not without further premises) arguments that the *assessor's* epistemic state ever needs to be relevant to determining the truth or falsity of an epistemic modal claim.

They're arguments, that is, that we need contents that take truth-values relative to something more than a world – since we've got the same contents being expressed by our various speakers in the examples, all of whom are worldmates, so that each of their contexts will supply the same world to the circumstance of evaluation that's relevant to the determination of utterance truth. But they're not arguments (at least, not without further premises) that particular utterances don't get their truth-values once and for all, or that we need to allow a role for contexts of assessment in the assignment of truth-values to sentences in context.

These sorts of common-content arguments are, in fact, most naturally understood as arguments that utterances with the same content can get different "once and for all" truth-values, *not* as arguments that we need to relativize utterance truth to contexts of assessment.

(For more on nonindexical contextualism, see MacFarlane 2009. For extensive discussion and criticism of these sorts of arguments for the view that propositional truth is relative to more than a world, see Cappelen and Hawthorne 2009.)

One argument, which is an elaboration of the disagreement argument above, *does* seem to support (if it's successful) not just stability of content, but also a role for context of assessment in the determination of truth-values for utterances. This is the argument from *faultless disagreement*.

There are two key moving parts of the argument from faultless disagreement: First, there's an imposition of a requirement that the relevant cases be classified by the correct theory as cases of *disagreement*. In the simplest arguments from faultless disagreement, this requirement is understood as a requirement that A's utterance semantically expresses the negation of B's utterance.

Second, there's a requirement of *faultlessness*. Both have to be getting it right, "from their own perspective". This isn't quite enough, however, as we've just seen that this by itself doesn't motivate a relativism of a kind that allows contexts of assessment to play a role in the determination of utterance truth. To accommodate this, all we need is a standard contextualist view (or a nonindexical contextualist view) that makes both A's and B's utterances come out to be true (true once-and-for-all).

But if we add a further requirement – we could call it a requirement of *strong faultless disagreement* – then we really do have an argument for a kind of relativism that allows utterance truth to be sensitive to context of assessment. To motivate this sort of relativism, we need to impose a requirement, not just that A is correct to regard his own utterance as true, but also that A is correct to regard B's utterance as *false*. (And v.v. for B.) This really does force us into a full-bloodedly relativist view, because then we have a pair of particular utterances – A's and B's – that need to take one truth-value for A, and the other for B. (Rather than just a pair of utterances that need to take different once-and-for-all truth-values, despite having the same content.)

(This sort of argument features prominently in relativist arguments on other topics, but less in arguments for relativism about epistemic modals. For some applications, see Kölbel 2002, 2008, 2009, Lasersohn 2005, 2009, Stephenson 2007, MacFarlane 2007, Egan forthcoming. For some opposition, see Sundell 2009 and von Fintel and Gillies 2008.)

This sort of argument from strong faultless disagreement is a member of a family of arguments meant to show that particular token utterances can be true relative to one assessor and false relative to another (within a world). That is, they're meant to show that utterances don't get truth-values once and for all, but only relative to a context of assessment. The argument from faultless disagreement bears a lot of the weight in discussions of relativism about other domains (in particular, about personal taste), but other members of the same family of arguments have standardly borne most of the weight in the case of epistemic modals.

One such argument (prominent in MacFarlane forthcoming) is based on *retraction* phenomena. There's a common conversational phenomenon of criticism and correction, in which one party to the conversation takes another to task for some previous utterance. (I say, "the cat is on the mat". You say, "nuh uh, the cat is not on the mat".) Abstracting away from a lot of detail, the standard responses to such criticism fall into two broad categories: *concession/retraction responses*, and *sticking to one's guns*.

Attempts at a certain sort of criticism and correction, which target utterances featuring context-sensitive vocabulary, after a relevant change of context has occurred in between the original utterance and the challenge, often evoke a distinctive kind of sticking-to-one's-guns: an *impatience/insistence* response, marked by such replies as, "oh, come on", "you know that's not what I meant", and "don't be a jerk".

For example:

The Drive

Dave: "It's hilly here"

(They drive along, out into the plains, changing the contextual features relevant to determining semantic value of “here”.)

Tim: “Nuh uh. It’s not hilly here at all.”

Dave: “Oh, come on. You know that’s not what I meant. Don’t be a jerk.”

Note that the alternative concessive/retracting response from Dave sounds weird and inappropriate here – it would be bizarre for Dave to reply to Tim’s challenge with any of the kinds of replies that mark this sort of response. (Such as, for example, “oh, I guess I was wrong”, “I take it back”, “my bad”, or “okay, scratch that, then”.) A similar pattern (I’ll say a bit more about just what the pattern is in a minute) seems to hold for context-sensitive vocabulary in general. It’s easy to construct parallel examples for “tall”, “nearby”, etc.

What’s happened here is that, between the utterance and the challenge, there’s been a shift in the relevant features of the conversational context. Prior to the shift, at the time of the original utterance, the context was one in which assertions of “it’s hilly here” would be true. After the shift, the context has changed in such a way that an assertion of “it’s hilly here” made in the new, post-change context would be false. More generally: the context has changed in such a way that, while the speaker ought to have taken himself to be in a position to truly assert what he did in the original context, the speaker ought no longer take himself to be in a position to truly assert the same sentence in the new, post-change context. Let’s call such a change in the conversational context an *undermining* context change, since such a change undermines the speaker’s ability to felicitously re-assert the same sentence. (The above is my best effort to state the phenomenon in terms that are neutral between contextualism and relativism. I fear that I may not quite have succeeded, but I hope I have done well enough to allow you to identify the phenomenon.)

With epistemic modals, on the other hand, we see a different pattern of appropriate reactions to challenges after undermining context changes:

Jane Emerges

Jim: "Bob might be in his office."

(Jane, previously not a participant in the conversation, introduces herself into the conversation, bearing a photograph of Bob in Amsterdam, just posted to Facebook seconds ago, changing the contextual features relevant to fixing the semantic value of "might".)

Jane: "Nuh uh. There's no way Bob is in his office – see, here's a picture of him going into a bar in Amsterdam just seconds ago."

Jim: "Oops, my bad – I take it back."

Note that the alternative, impatience/insistence response seems out of place here. It would be off for Jim to respond to Jane's challenge by saying, e.g., "oh, come on – you know that's not what I meant".

Retraction arguments against contextualism (and in favor of relativism) about epistemic modals are based on this pattern of response to challenges and criticisms after undermining context changes, which don't look like the ones that a contextualist theory would predict.

The most interesting use of these examples, for our purposes, is as an argument that a single, dated utterance can get different truth-values relative to different contexts of assessment: Assessed from its original context of utterance, Jim's claim is true. Assessed from the post-shift context of assessment, the very same utterance is false. That's why the speaker *retracts*, rather than just refraining from uttering the same sentence in the new, changed context.

The challenge to the contextualist is to explain the retraction phenomena while insisting that the initial utterance retains the same truth-value relative both to the original context of utterance and the later context of retraction. The relativist argues that this is going to be impossible to do.

If the contextualist says that the original utterance was (and remains) *true*, it's difficult to explain why the speaker retracts in the face of the later challenge. If the contextualist says that the original utterance was (and remains) *false*, it looks as if she'll have to say that that's because Jane's information was relevant all along. And this, according to the relativist, is going to make the truth-conditions for

“might” claims too demanding. Once we allow Jane to be (and to have been all along) a member of the group whose evidence makes a difference to the truth of Jim’s claim, it’s hard to identify a principled place to stop short of *everybody who ever thinks about the utterance* being in the relevant group. And if that’s so, “might” claims are going to look incredibly risky, and we’ll barely ever be in a position to felicitously assert them, since it will be so easy for them to be false.

Responses to this sort of argument from defenders of contextualism fall into two categories: challenges to the data, and attempts to accommodate the phenomenon within a contextualist theory. (Often, these moves complement each other – the contextualist offers a more sophisticated theory that accommodates a lot of the troublemaking phenomena, and then challenges the data to argue that the remaining un-accommodated phenomena aren’t genuine, or aren’t genuinely problematic.)

One compelling response of the data-challenging type is to point out that the tendency to retract isn’t so universal as the initial versions of the relativist argument suggest. We don’t *always* get retraction in the face of challenges after an undermining context change. We sometimes get insistence and impatience. Here is a case from von Fintel and Gillies (2008):

The Keys

Alex: The keys might be in the drawer.

Billy: (Looks in the drawer, agitated.) They’re not. Why did you say that?

Alex: Look, I didn’t say they *were* in the drawer. I said they *might* be there—and they might have been. Sheesh.

(von Fintel and Gillies 2008, p.81)

Alex’s response seems completely appropriate. He should, in this case, dig in his heels, refuse to retract, and insist on the correctness of his past “might” claim. So it looks as if the phenomena here aren’t as the standard relativist arguments portray them – at least, they’re not as the simple versions of the relativist arguments portray them. There are two ways to use this against the relativist.

One use of insistence phenomena is *defensive*: Since the data aren't as the relativist described them, the arguments from retraction phenomena fail, and the world is still safe for contextualism after all. While it certainly does seem to be true that the data aren't as simple as the standard relativist arguments make them out to be, it's not clear how successful this reply is. Plausible versions of this response will still leave us with *some* cases of retraction after what looks like a relevant context change. (In particular, after the introduction of new participants to the conversation.) These retractions will still need explaining, and they aren't predicted by at least the first sorts of contextualist theories that spring to mind. So this can't be a *complete* defense – there's also a need for an elaboration on the contextualist theory that allows it to accommodate the remaining cases of genuine retraction. (One such theory is offered by von Fintel and Gillies forthcoming.)

Another use of insistence phenomena is *offensive*: The insistence data are actually quite uncomfortable for the relativist. Relativist theories seem to predict that assessors' truth-value assessments will *always* be based on whether the utterance under evaluation is true relative to their own present context of assessment (or relative to the pair of the context of utterance and their own present context of assessment). They predict, therefore, that we'll *never* see the kinds of insistence on past correctness that we see in *The Keys*.

The relativist needs to either explain these away or provide a story that accounts for them. One way to offer such an account is to follow Peter Lasnik's (2005, 2009) lead, and allow for two different sorts of truth-value attributions: one that tracks truth relative to the assessor's present context, and one that tracks truth relative to the speaker's context of utterance. In the cases of retraction, what's at issue is whether original utterance is true in the conversationalists' present contexts of assessment. Correctly recognizing that it's false relative to his own present context of assessment, the speaker retracts. In the cases of insistence, what's at issue is whether the original utterance was true in the original context of utterance. Correctly recognizing that it is, the speaker insists.

Relativists of both the *de se* and MacFarlanian types have the resources available to make this distinction. (Though in MacFarlane's framework it will be

phrased slightly differently – as the distinction between truth at the pair $\langle C_U, C_A \rangle$ and truth at the pair $\langle C_U, C_U \rangle$ - and MacFarlane himself seems not inclined to adopt it.) Both also face the challenge of explaining, in a way that's not *ad hoc*, just why the one kind of evaluation is relevant in the retraction-generating conversations, and the other kind is relevant in the insistence-generating conversations.

Another sort of argument for variations in the truth-value of particular utterances is what are often known as *eavesdropper* arguments, based on the assessments of utterances by third parties who are not participants in the conversation in which the utterance takes place. Unlike what we find with standard context-sensitive expressions, and unlike what one would expect from a context-sensitive expression, evaluators who are not parties to the conversation in which an utterance occurs systematically base their assessments of utterances of epistemic modal claims on (their views about) their *own* evidential state, not on (their views about) the evidential state of the speaker, speaker's group, etc.

For example, here is a contrasting pair of cases from MacFarlane (forthcoming):

George and Sally

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?

The hope is that you will have answered "no" to the first question and "yes" to the second. And this does indeed seem to be the reaction most people have. This example, as MacFarlane deploys it, is meant to do a number of things: It's supposed to identify an important difference between, e.g., "Joe might be in Boston" and, e.g.,

“As far as I know, Joe is in Boston”, and thereby head off any equivalence claims that we might otherwise have been tempted to make. More importantly, it’s also supposed to show that our assessments of the truth-values of other people’s “might” claims, even those that take place in conversations that we’re not parties to, are based on what we think about *our own* epistemic state, not on what we think about the epistemic state of the speaker, or the people who are party to the conversation in which the utterance takes place.

Here is another way of getting at the same phenomenon: Moore (1962) notes that one of the markers of the fact that “It’s possible that Hitler is now dead”, and “Hitler may be dead” are *epistemic* possibility claims is that one can use “I know that he’s not” to deny them. This is an instance of a striking general feature of epistemic possibility claims, which is illustrated above, and also by Hawthorne’s inspirational footnote: *anybody* can use first-person knowledge claims to deny them – not just addressees, and not just participants in the conversation in the course of which the claim is made.

The relativist argument here is that contextualists can’t give an adequate explanation of this fact, while a relativist view is ideally suited to explain it. Contextualism, the relativist charges, predicts that competent extra-conversational assessors will defer, in their assessments of utterances for truth and falsity, to the context of utterance for content-determination and index-selection. So they shouldn’t, in general, take their own epistemic state to be relevant, any more than they should take their own *location* to be relevant when they’re assessing occurrences of “here” claims in conversations that they’re not participants in.

One candidate contextualist reply here is to say the extra-conversational evaluator’s epistemic state really *was* relevant all along, and so the utterance is just plain false. The trouble with this is that it’s difficult to tell such a story without making the truth-conditions for epistemic “might” claims excessively demanding.

We don’t, as a matter of fact, need to rule out the presence of better-informed lurkers in the closet in order to take ourselves to be in a good position to make “might” claims. And even if we *strongly suspect* lurkers – even if we’re *certain* there’s a lurker – we still don’t take our “might” claims to be hostage to the lurker’s

knowledge. An example to make this clear: Tony Soprano can felicitously say to Paulie, “be careful, the Feds might have a guy outside your house”, even if they’re sure that there’s an FBI agent listening in who knows for sure whether there’s a guy outside Paulie’s house. This is so even if Tony is pretty confident that there *isn’t* a Fed outside Paulie’s house, and that the eavesdropper knows this. And it’s so even if he thinks the eavesdropper is hiding in the closet rather than listening in on a wiretap. (Tony could also say, “there might be a guy outside the house, and there might not. But it’s best to be careful”. At least one conjunct of this is sure to be false if the eavesdropper’s knowledge is relevant.)

The relativist proposes to explain these phenomena by offering a semantic theory according to which a single, dated utterance of an epistemic modal claim can take different truth-values relative to different assessors, so that it’s true for (e.g.) the speaker, and false for the extra-conversational assessor.

Here too, there are two types of reply available to the contextualist: They can challenge the data, or they can offer a way for a contextualist theory to accommodate the phenomena. (As well as the mixed strategy, of challenging the data to cut down the range of phenomena to be accommodated, and then telling a story that accommodates the remainder.)

Many of the standard presentations of eavesdropper arguments (see for example Egan, Hawthorne and Weatherson 2005 and Egan 2007) rely on intuitions about assessors’ utterances of “that’s true” or “that’s false” in response to overheard epistemic modal claims. One concern about such intuitions is that just what they show depends on how some difficult issues about propositional anaphora wind up playing out. In particular, one way to resist the relevance of those intuitions is to offer alternative hypotheses about the reference of “that” in the relevant utterances, such that what the assessor is delivering an assessment of isn’t the original utterance, or the proposition it expressed. (See, for example, von Fintel and Gillies 2008 for some such replies.)

There are a couple of strategies available to the relativist in order to counter this reply. One is to tailor the responses in the examples in order to make the anaphora unambiguous. (Narrowing in, for example, on *Bob might be in his office*

rather than the prejacent *Bob is in his office* as the target of assessment by noting that it seems in order for the evaluator to say, e.g., “that’s false, there’s no way Bob is in his office”, or “that’s false, it’s impossible that Bob is in his office”.) Another alternative is to restate the argument in anaphora-free terms, as I have followed MacFarlane in doing here.

Besides challenging the relativist’s claims about the data regarding extra-conversational assessments, another possible strategy is to adjust the contextualist theory in order to accommodate them. One option here is to tighten up the truth-conditions for epistemic “might” claims, so that the speakers *are* speaking falsely, but blamelessly. The intra-conversational assessors are making the same sort of blameless mistake as the speaker, and the extra-conversational assessors are just getting their truth-value assessments right. The most natural way to do this is to say that eavesdroppers are, in general, to be included in the relevant group whose epistemic state makes a difference to the truth or falsity of epistemic modal claims. This strategy is subject to the same kind of concern as the analogous response to retraction arguments above: It’s very hard to do this without making the truth-conditions for epistemic modal sentences (a) so demanding that we’re basically never in a position to assert them, and (b) demanding in ways that make bad predictions about when we should take ourselves to be in a position to assert them. (Though see von Fintel and Gillies forthcoming for a contextualist strategy that seeks to avoid this problem)

There is, obviously, much more to say about all of this. But I hope that this section has succeeded in displaying at least the standard motivations for relativism about epistemic modals, and the first step or two of the ensuing debates.

Conclusion

I’ve tried in the preceding to lay out, in broad strokes, the candidate relativist views about epistemic modals, and the contextualist orthodoxy that they’re seeking to replace. I’ve also tried to lay out, in similarly broad strokes, the main arguments relativists offer for their views, and the first few steps of the ensuing debate between the relativists and the anti-relativists about whether or not those

arguments succeed. While I've looked at some of the anti-relativists' responses to the relativists' positive arguments, one thing that I have not done, do to limitations of time and space, is survey the various positive arguments against relativism that its opponents offer. I will close, then, with some suggestions, for those who are interested, of points of entry into that area of the literature:

For some positive arguments against (or at least doubts and concerns about) relativism in general, and not specifically about epistemic modals, see for example Wright 2008, Cappelen and Hawthorne 2009, Garcia-Carpintero 2008, Evans 1985, and von Fintel and Gillies 2008.

For some arguments against relativism about epistemic modals in particular, see for example Dietz 2008, von Fintel and Gillies forthcoming, Schaffer forthcoming, Yalcin 2007 and forthcoming, and Swanson forthcoming.

References:

- Bach, Kent (forthcoming) Perspectives on Possibilities: Contextualism, Relativism, or What?, forthcoming in *Epistemic Modality*, Andy Egan and Brian Weatherson, eds. Oxford University Press.
- Cappelen, Herman (2008a) Content Relativism and Semantic Blindness, in *Relative Truth*, Manuel Garcia-Carpintero and Max Kolbel, eds. Oxford University Press, pp. 265-86.
- Cappelen, Herman (2008b) The Creative Interpreter, *Philosophical Perspectives* 22, pp. 23-46.
- Cappelen, Herman and Hawthorne, John (2009) *Relativism and Monadic Truth*, Oxford University Press.
- Cappelen, Herman and Lepore, Ernest (2005) *Insensitive Semantics*, Wiley-Blackwell.
- Chierchia, Gennaro. 1989. Anaphora and attitudes de se. In *Semantics and contextual expression*, ed. Renate Bartsch, Johan van Benthem and Peter van Emde Boas. Dordrecht: Foris.
- DeRose, Keith (1991) Epistemic Possibilities, *Philosophical Review* 100:4, 581-605.
- Dever, Josh (forthcoming) Epistemic Modals, in *Routledge Companion to Epistemology*, Duncan Pritchard, ed., Routledge.
- Dietz, Richard (2008) Epistemic Modals and Correct Disagreement, in *Relative Truth*, Manuel Garcia-Carpintero and Max Kolbel, eds. Oxford University Press, pp. 239-62.
- Egan, Andy, John Hawthorne and Brian Weatherson, Epistemic Modals in Context, in *Contextualism in Philosophy*, Gerhard Preyer and Georg Peter, eds., Oxford University Press. pp. 131-170.

- Egan, Andy (2007) Epistemic Modals, Relativism and Assertion, *Philosophical Studies* 133, pp.1-22.
- Egan, Andy (2009) Billboards, Bombs and Shotgun Weddings, *Synthese* 166, pp. 251-279.
- Egan, Andy (forthcoming) Disputing about Taste, forthcoming in *Disagreement*, Richard Feldman & Fritz Warfield, eds., Oxford University Press.
- Evans, Gareth (1985) Does Tense Logic Rest on a Mistake?, in his *Collected Papers*, Clarendon Press, pp.342-63.
- von Fintel, Kai and Gillies, Anthony (2008) CIA Leaks, *Philosophical Review* 117:1, pp. 77-98.
- von Fintel, Kai and Gillies, Anthony (forthcoming) 'Might' Made Right, forthcoming in *Epistemic Modality*, Andy Egan and Brian Weatherson, eds. Oxford University Press.
- Garcia-Carpintero, Manuel (2008) Relativism, Vagueness and What is Said, in *Relative Truth*, Manuel Garcia-Carpintero and Max Kölbel, eds. Oxford University Press, pp. 129-154.
- Hacking, Ian (1967) Possibility, *Philosophical Review* 76:2, pp143-68.
- Hawthorne, John (2006) *Knowledge and Lotteries*, Oxford University Press.
- Kaplan, David (1989) Demonstratives, in Joseph Almog, John Perry and Howard Wettstein, eds., *Themes from Kaplan*. Oxford University Press.
- Kölbel, Max (2002) *Truth Without Objectivity*, Routledge.
- Kölbel, Max (2003) Faultless Disagreement, *Proceedings of the Aristotelian Society* 104:1, pp53-73.
- Kölbel, Max (2009) The Evidence for Relativism, *Synthese* 166, pp375-95.
- Kratzer, Angelika (1977) What 'Must' and 'Can' Must and Can Mean, *Linguistics and Philosophy* 1, pp. 337-355.
- Kratzer, Angelika (1981) The Notional Category of Modality. In *Words, Worlds, and Contexts: New Approaches to Word Semantics*, ed. H. J. Eikmeyer and H. Rieser, 38-74. Berlin: de Gruyter. Reprinted in *Formal Semantics: The Essential Readings*, ed. P. Portner and B. H. Partee, 2002. Oxford: Blackwell.
- Kratzer, Angelika (1986) Conditionals, *Chicago Linguistics Society* 22, pp.1-15
- Kratzer, Angelika 1991. "Modality". In *Semantics: An International Handbook of Contemporary Research*, edited by Arnim von Stechow & Dieter Wunderlich, pp. 639-650. Berlin: de Gruyter.
- Laserson, Peter (2005) Context Dependence, Disagreement, and Predicates of Personal Taste, *Linguistics and Philosophy* 28, pp643-86.
- Laserson, Peter (2009) Relative Truth, Speaker Commitment, and Control of Implicit Arguments, *Synthese* 166, pp359-74.
- Lewis, David (1980). Index, Context and Content. In *Philosophy and Grammar*, Stig Kanger and Sven Ohman, eds., D. Riedel. Reprinted in David Lewis, *Papers in Philosophical Logic*, Cambridge University Press, pp. 21-44.
- MacFarlane, John (2005) Making Sense of Relative Truth. *Proceedings of the Aristotelian Society* 105, pp. 321-39.
- MacFarlane, John (2007) Relativism and Disagreement. *Philosophical Studies* 132 pp. 17-31.
- MacFarlane, John (2009) Non-Indexical Contextualism, *Synthese* 166, pp.231-50.

- MacFarlane, John (forthcoming) Epistemic Modals are Assessment-Sensitive, forthcoming in *Epistemic Modality*, Andy Egan and Brian Weatherson, eds. Oxford University Press.
- Moore, George Edward (1962) *Commonplace Book*, George Allen & Unwin.
- Predelli, S. (1996). Never put off until tomorrow what you can do today. *Analysis*, 56(2), 85–91.
- Predelli, S. (1998a). I am not here now. *Analysis*, 58(2), 107–115.
- Predelli, S. (1998b). Utterance, interpretation and the logic of indexicals. *Mind & Language*, 13(3), 400–414.
- Richard, Mark (2004) Contextualism and Relativism, *Philosophical Studies* 119, pp214-42.
- Richard, Mark (2009) *When Truth Gives Out*, Oxford University Press.
- Schaffer, Jonathan (forthcoming) Contextualism for Taste Claims and Epistemic Modals, forthcoming in *Epistemic Modality*, Andy Egan and Brian Weatherson, eds. Oxford University Press.
- Stephenson, Tamina (2007) Judge Dependence, Epistemic Modals, and Predicates of Personal Taste, *Linguistics and Philosophy* 30, 487-525.
- Sundell, Timothy (2009) *Conflict and Content*, PhD dissertation, University of Michigan.
- Swanson, Eric (forthcoming) How Not to Theorize about the Language of Subjective Uncertainty, forthcoming in *Epistemic Modality*, Andy Egan and Brian Weatherson, eds. Oxford University Press.
- Teller, Paul (1972) Epistemic Possibility, *Philosophia* 2:4, pp303-20.
- Wright, Crispin (2008) Relativism about Truth Itself: Haphazard Thoughts about the Very Idea, in *Relative Truth*, Manuel Garcia-Carpintero and Max Kolbel, eds. Oxford University Press, pp. 157-186.
- Yalcin, Seth (2007) Epistemic Modals, *Mind* 116, pp. 983-1026
- Yalcin, Seth (forthcoming) Non-Factualism about Epistemic Modality, forthcoming in *Epistemic Modality*, Andy Egan and Brian Weatherson, eds. Oxford University Press.