

# 9

## EPISTEMIC MODALS

SETH YALCIN

*Massachusetts Institute of Technology*

### 1

I have a puzzle about epistemic modals.<sup>1</sup> I spell it out in the next two sections (§2-3). The puzzle teaches, I eventually say, that ‘possibly’, ‘probably’, and the rest of the epistemic modals do not designate functions which yield, as output, propositions. As a result, the puzzle tells against standard truth-conditional approaches to the semantics of epistemic modals (§4), and also against more recent assessment-sensitive approaches to their semantics (§5-6). I spend the balance of the paper trying to develop the view that epistemic modals modulate force, not content.

### 2 A puzzle about epistemic modals

Famously, sentences of the form ‘ $\phi$  and I don’t know  $\phi$ ’, such as

- (1) It isn’t raining and I don’t know it isn’t raining.

are odd, contradictory-sounding, and generally unassertable at a context.<sup>2</sup> This is taken to be *prima facie* puzzling, for such sentences might well

---

<sup>1</sup> By ‘epistemic modals’ I mean operators or modal verbs like ‘possibly’, ‘it might be that’, ‘it must be that’, ‘probably’, ‘it’s likely that’, ‘it must be that’, etc., interpreted under their “epistemic” readings—i.e., on the readings that somehow convey epistemic information. (The semantic challenge, obviously, is to spell out the “somehow”.)

<sup>2</sup>The same effect is achieved if ‘know’ is replaced by ‘believe’, ‘suspect’, or any other of related doxastic or epistemic attitude verbs, as is well known. I am sticking to ‘know’ here for reasons that will emerge.

express truths. Call sentences that are odd in this way *Moore-paradoxical*. The challenge—or one challenge—that Moore-paradoxical sentences are supposed to present is that of explaining their oddness, given no straightforward appeal to classical contradiction is available.

Now it is thought that, just like the certain attitude verbs, epistemic modals can be used to generate Moore-paradoxical sentences.<sup>3</sup> For sentences like

- (2) It isn't raining and it might be raining.
- (3) It isn't raining and it is possibly raining.

are odd, contradictory-sounding, and generally unassertable at a context, too; and the modals they contain are epistemic, conveying, like certain attitude verbs, some sort of information about a knowledge state or set of states. Further, just as with canonically Moore-paradoxical sentences, it isn't plausible to maintain that (2) and (3) really express logical contradictions. For this would require MIGHT  $\phi$  and POSSIBLY  $\phi$  to each logically entail  $\phi$ ; and that would be absurd, because it would make MIGHT and POSSIBLY factives. So a natural conjecture concerning these sentences, then, is that we have the same basic puzzle—Moore's paradox—in a new guise. And indeed, standard approaches to the semantics of epistemic modals tend to conform to this thought, if only implicitly. They conform to it by having 'It might be raining' entail 'I don't know it's not raining' relative to a context and speaker, so that (2) entails (1). The oddity of (2) can then be blamed on this entailment. (Repeat the story for (3)—POSSIBLY  $\phi$  is assumed to be equivalent to MIGHT  $\phi$  on the epistemic readings of these modals.) The thought is that these sentences are odd because they entail garden-variety Moore-paradoxical sentences; hence their oddity will be explained by whatever explains garden-variety Moore-paradoxicality. The standard explanations of the garden variety cases are pragmatic in form, pointing to some kind of conflict between the content of sentences like (1) and the norms which govern assertion.

I think this misses something. Agreed that (1), (2), and (3) are one and all odd, noncontradictory, contradictory-sounding, and unassertable. Still, there is a nontrivial difference between classically Moore-paradoxical sentences like (1), on the one hand, and epistemically modalized counterparts of these sentences, like (2) or (3). The difference is that there are some propositional attitudes we can coherently take towards the content

<sup>3</sup>The thought occurs explicitly in, e.g., Hacking (Hacking, 1967, p.150) and Stalnaker (Stalnaker, 1981, p.99-100).

of sentences like (1), but not towards the content of sentences like (2) or (3).

For instance, we can coherently *suppose* (1), but not (2) or (3).<sup>4</sup> That is, there is an attitude, namely supposition, that we can take towards what (1) says, but not what (2) or (3) say. I can easily suppose that it isn't raining and that I don't know it—there is no difficulty at all here. (That we can easily suppose such things is partly why, after all, we don't think Moore-paradoxical sentences express logical contradictions.) In contrast, it is just not possible to coherently suppose that it isn't raining and it might be raining, or that it isn't raining and it is possibly raining. Asked to suppose one of these, one hardly understands the question.

Other, related attitudes yield the same asymmetry. Take *conceiving*, or *conceiving it to be the case that*: I can conceive (both) that it isn't raining and that I don't know it, but not (both) that it isn't raining and that it might be. To vary the language, I cannot conceive of a situation which is truly described by both of the conjuncts in (2). *Imagining* and *entertaining* are two more attitudes yielding the same asymmetry. Perhaps all these attitudes form a class, with some other interesting shared properties. Perhaps these attitudes are really just slight variations on one another. Whatever; the point I want to make is simply that there are some attitudes that we can take towards what sentences of the form  $\ulcorner \neg\phi$  and I don't know  $\neg\phi \urcorner$  say, but not towards what sentences of the form  $\ulcorner \neg\phi$  and MIGHT  $\phi \urcorner$  or  $\ulcorner \neg\phi$  and POSSIBLY  $\phi \urcorner$  say. (1), (2), and (3) are indeed one and all odd, noncontradictory, contradictory-sounding, and unassertable; but (2) and (3) are, in addition, un-supposable (inconceivable...). This additional feature, I hope you will agree, is surprising and puzzling. No alleged entailment relation between (2)/(3) and (1) will explain this fact, for (1) is itself perfectly supposable. What is especially puzzling about (2) and (3), I think, is that they are each *both* un-supposable *and* noncontradictory. The presence of one of these properties is typically good evidence for the absence of the other—especially in conjunctions neither of whose conjuncts are themselves contradictory. To be sure, there do exist other sentences with both of these properties. But the clearest examples are ones which are un-supposable simply because they are computationally too difficult to process, and (2) and (3) are too unremarkable in syntactic or conceptual complexity for this to be a plausible explanation. So the puzzle is this: given that (2) and (3) are both noncontradictory, and given that both are unremarkable in syntactic or conceptual complexity, why can't we suppose whatever it is these sentences say?

<sup>4</sup>I am using 'suppose' in the ordinary English sense, not the very thin, technical, quasi-syntactic notion of *supposing* involved when we speak of, e.g., supposing for *reductio*.

So far, in motivating this puzzle, what I have mainly helped myself to are intuitions about truth-conditions. The thought is just this: if either of (2) and (3) had (non-null) truth-conditions, one would be able to consider (suppose, conceive) one of those possibilities meeting the conditions they specify, whatever they are. One would be able to do this in just the way we can do it for the proposition expressed by (1) (again, relative to a context and speaker). The claim that you can't do this is thus a kind of judgment about truth-conditions: it is the judgment that (2) and (3) are true under no conditions. Combined with the independently motivated assumption that neither (2) nor (3) are contradictory, one is left with a puzzle about what, if anything, the truth-conditions of these conjunctions could possibly be. Ultimately, this is a puzzle about the semantics of epistemic modals: what must their semantics be, if these modals can give rise to sentences with the strange properties of (2) and (3)?

### 3 Embedding

I will broach this question shortly, but let me say a little more, first, about the puzzle. As I said, I have motivated it by appeal to intuitions about truth-conditions. But that there is a puzzle here need not be left to rest wholly on intuitions about truth-conditions. The puzzle also manifests itself in judgments about the acceptability or interpretability of sentences in context. To see this, we must look beyond (1), (2), and (3), for these sentences are uniformly unacceptable; they all make for lousy conversational moves. To get an asymmetry in acceptability between (2)/(3) and (1), we must try embedding these clauses. And indeed, when we try this, we do find that canonically Moore-paradoxical sentences are easier to embed than their epistemically modalized counterparts.<sup>5</sup>

I will give two sample embedding environments where the asymmetry arises. The first environment will be obvious:

(4) Suppose that it isn't raining and I/you/we don't know it isn't raining.

(5) # Suppose that it isn't raining and it might be raining.

Here we have a clear contrast in acceptability. When 'suppose' embeds a Moore-paradoxical sentence as in (4), the result is acceptable and perfectly

---

<sup>5</sup>To avoid excessive repetition, I will restrict myself to a comparison of (1) and (2), but everything that follows can be repeated exactly with (1) and (3).

interpretable. When ‘suppose’ embeds the epistemically modalized counterpart of such a sentence as in (5), however, the result is defective and not coherently interpretable. To generalize:

- (6) # Suppose [  $\neg\phi$  and MIGHT  $\phi$  ].

The unacceptability judgment we have with (5) corroborates the thought voiced above, viz., that what the embedded clause expresses is not supposable. For if it were supposable, this imperative would presumably have coherent satisfaction conditions, and hence be coherently interpretable.

Here is the second embedding case. Moore-paradoxical sentences can occur happily in the antecedent position of an indicative conditional, whereas their epistemically modalized counterparts cannot:

- (7) If it isn’t raining and I don’t know it isn’t raining, then (well, still) it isn’t raining.
- (8) # If it isn’t raining and it might be raining, then (well, still) it isn’t raining.

Again we have a contrast. True, (7) flouts a Gricean maxim or two<sup>6</sup>, but it is plainly not odd in anything like the Moorean sense. Interpretation is relatively straightforward—what (7) says seems to be trivially true. The conditional (8), on the other hand, is incoherent and uninterpretable. It is worth emphasizing that the trouble with (8) is its antecedent, and not some interaction effect between antecedent and consequent. While there are a variety of acceptable ways to finish off a conditional containing (1) in antecedent position, as witnessed by (7) above and, perhaps more natural, by (9) below:

- (9) If it isn’t raining and I don’t know it, then there is a truth I don’t know.

there is just no happy way to finish off a conditional containing (2) in antecedent position:

- (10) # If it isn’t raining and it might be raining, then ...

Such conditionals appear to be, as it were, already hopeless. Generalizing:

---

<sup>6</sup>Viz., relevance (what the conditional says seems to be trivially true) and the second maxim of quantity (given the usual presupposition in a discourse context that the antecedent of an indicative is not known, the second conjunct of the antecedent may sound superfluous).

(11) # If  $\neg\phi$  and MIGHT  $\phi$ , then  $\psi$ .

It is as if the oddity of sentences like (2) somehow “projects” out of embedded contexts, infecting the larger constructions that contain them.<sup>7</sup>

A *prima facie* worry about the foregoing embedding considerations needs to be parried. I am suggesting that (2) is to blame for the unacceptability of (5) and (8), that it is the clause making these larger constructions defective. But it may be replied that the real trouble with (5) and (8) is that they both embed epistemic modals, and that embedding epistemic modals—at least in these environments—just isn’t allowed, perhaps for (yet to be detailed) syntactic reasons. According to this objection, the conjunction (2) is a red herring, and the more telling generalizations are these:

(12) # Suppose MIGHT  $\phi$ .

(13) # If MIGHT  $\phi$ , then  $\psi$ .

Now this worry is not unmotivated, for it is known that epistemic modals are not freely embeddable.<sup>8</sup> Nevertheless, (12) and (13) do not appear to be quite correct. It is true that some do find

(14) ? Suppose it might be raining.

marked. But most appear to find it acceptable, and (in particular) far better than (5), which is uniformly rejected. So the judgments about (14) do not appear to be robust enough to account for the defectiveness of (6). As for conditionals with epistemic modal antecedents, such constructions are indeed uncommon. But they do undoubtedly exist:

(15) If the restaurant might be closed, we should call ahead.

(16) If Smith might be in his office, then Jones might be in her office.

(17) If Dan probably won’t be at the party, then I’m going to just stay home.

---

<sup>7</sup>As with embedding under ‘suppose’, one may find it tempting to trace the uninterpretability of indicative conditionals with the form of (11) to some more basic fact about the incoherence of supposition-like attitudes towards what sentences of the form ‘ $\neg\phi$  and MIGHT  $\phi$ ’ say. One familiar idea is that the interpretation of indicative conditionals essentially requires something like supposition of the antecedent. (See, for instance, Edgington (1986) and references cited therein.) Such a view is supported by the facts so far presented, for it predicts, correctly, that our un-supposable ‘ $\neg\phi$  and MIGHT  $\phi$ ’-sentences should make for defective antecedents.

<sup>8</sup>See, e.g., von Stechow and Iatridou (2003) and references cited therein.

So (13) could not be right. It seems reasonable to assume, then, that (6) and (11), repeated:

(6) # Suppose [  $\neg\phi$  and MIGHT  $\phi$  ].

(11) # If  $\neg\phi$  and MIGHT  $\phi$ , then  $\psi$ .

capture the correct level of generalization.

To summarize, we have two related differences between paradigmatically Moore-paradoxical sentences and their epistemic modal counterparts. Paradigmatically Moore-paradoxical sentences are unassertable, but supposable. Their epistemic modal counterparts, in contrast, are both unassertable and unsupposable. Paradigmatically Moore-paradoxical sentences are coherently embeddable (e.g.) under ‘suppose’, and in the antecedent position of a conditional. Not so with their epistemic modal counterparts. When these are embedded under ‘suppose’ or in the antecedent position of a conditional, the constructions that result are defective and uninterpretable. I will assume, as seems reasonable, that these two differences are facets of the same puzzle.

So far I have emphasized ‘might’ and ‘possibly’. But other operators are standardly classified as epistemic modals, in particular ‘probably’ and ‘likely’. It is time to point out that these operators give rise to exactly the same kind of effects. Thus these sentences are not supposable:

(18) It isn’t raining and it is probably raining.

(19) It isn’t raining and it is likely raining.

and the relevant embeddings are not acceptable:

(20) # Suppose it isn’t raining and it is probably raining.

(21) # Suppose it isn’t raining and it is likely raining.

(22) # If it isn’t raining and it is probably raining, then (still) it isn’t raining.

(23) # If it isn’t raining and it is likely raining, then (still) it isn’t raining.

Note also that the addition of modifiers to these operators (such as ‘very’ and ‘quite’) makes no difference.

What of epistemic ‘must’, the canonical dual of ‘might’? It is easy to verify that epistemic ‘must’ does generate all the effects I have described:

- (24) It isn't raining and it must be raining.
- (25) # Suppose it isn't raining and it must be raining.
- (26) # If it isn't raining and it must be raining, then (still) it isn't raining.

With epistemic 'must', however, a *prima facie* natural explanation of these effects suggests itself, one that does not arise for the other epistemic modals so far mentioned. The explanation is that 'must' is factive, making (24) literally contradictory, hence not conceivable, hence not sensibly embeddable in the above environments. Because that sounds plausible to me, I will not claim to have raised a puzzle about 'must'. Consequently, this modal will take a backseat in the discussion to follow. Only those who deny that 'must' is factive have a puzzle here.

So sentences like  $\lceil \neg\phi$  and MIGHT  $\phi \rceil$ ,  $\lceil \neg\phi$  and POSSIBLY  $\phi \rceil$ , and  $\lceil \neg\phi$  and PROBABLY  $\phi \rceil$  have their own proprietary weirdness. The weirdness needs explaining. What must the meanings of the epistemic modals be, if the facts are as I have just described?

## 4 Standard truth-conditional semantics for epistemic modals

It will be easier to start negatively, with what the meanings of the epistemic modals could not be, given the facts described. In this section, I will argue that the modals could not designate functions which yield, as output, propositions (understood as functions from worlds to truth-values).

I will focus on the standard versions of this view about the semantics of epistemic modals. According to standard approaches, epistemic modal statements (MIGHT  $\phi$ , MUST  $\phi$ , POSSIBLY  $\phi$ , PROBABLY  $\phi$ , etc.) have truth-conditions relative to a context (characterizable by a proposition) and amount, basically, to descriptions of epistemic states: their truth values turn on whether some specified epistemic state(s) have some specified property. A number of non-equivalent analyses exist falling under this rough characterization. My strategy will be to begin with a very simple version of this approach, saying why it fails to accommodate (predict) the weird properties of epistemic modals described above. Then I will try to generalize this basic objection, arguing that it survives even when the simple version of the standard approach is substantially modified along the two most obvious parameters. My focus will be entirely on epistemic 'might' and 'possibly' here; problems exactly analogous to the ones I will raise occur for other epistemic modals as well (e.g., 'probably'), but to discuss the

latter would introduce complications unnecessary to see the main thrust of the objection.

According to the simple analysis I want to start with—call it the *standard speaker-relative analysis*—epistemic modal statements express propositions about the speaker’s epistemic state.<sup>9</sup> In particular, when one says

It might be raining

one says that it is possible, for all one knows, that it is raining. Equivalently: one says that its raining is compatible with everything one knows. Equivalently again: what one says is true just in case one’s knowledge does not rule out the possibility that it is raining. Equivalently one last time: what one says is true just in case one does not know that it is not raining. The standard way to implement this idea semantically is to model ‘might’ in the metalanguage as an (existential) quantifier over worlds, the worlds compatible with what the speaker knows at the context of utterance. Reading  $\llbracket \alpha \rrbracket^c$  as “the semantic value of sentence  $\alpha$  relative to context of utterance  $c$ ”, we have

$$\llbracket \text{MIGHT } \phi \rrbracket^c \text{ is true at } w \text{ iff } \exists w' \in R_c(w) : \llbracket \phi \rrbracket^c \text{ is true at } w'$$

where the value of the restricted set of worlds quantified over is

$$R_c(w) = \text{the set of worlds compatible with what the speaker of } c \text{ knows at } w$$

(Same story for ‘possibly’.) Plainly, the particular set of worlds ‘might’ quantifies over is context-sensitive—it depends on the speaker and what she knows at the evaluation world  $w$ .<sup>10</sup> Relative to a context, the truth-conditions of ‘It might be raining’ are, on this story, equivalent to ‘I don’t know it isn’t raining’. The sentence thus describes the speaker’s epistemic state: what an utterance of MIGHT  $\phi$  or POSSIBLY  $\phi$  tells you is whether  $\phi$  is compatible with what the speaker knows.

That is the account. It is known to be inadequate.<sup>11</sup> But without rehearsing the usual objections, I hope it is already obvious that, given our previous observations, this semantics could not be correct. On this

<sup>9</sup>The idea here occurs to Moore (Moore, 1962, p.184) and Stalnaker (1981), Stalnaker (1984).

<sup>10</sup>Unless shifted by higher operators, the evaluation world is just the world of the context of utterance.

<sup>11</sup>See, for instance, the objections in Hacking (1967), DeRose (1991) and MacFarlane (2003).

account, ‘it might be raining’ and ‘I don’t know it isn’t raining’ have the same truth-conditions. Compositionally, they should make the same truth-conditional semantic contributions to the larger constructions in which they occur. Consequently, these two sentences ought to have the same truth-conditions:

- (1) It isn’t raining and I don’t know it isn’t raining.
- (2) It isn’t raining and it might be raining.

The semantic content of these two sentences, expressed relative to the same context, should therefore be true at all the same worlds. Moreover, this should be obvious to all but the semantically incompetent. But that is wrong: relative to a context, the worlds where (1) is true are simply not the worlds where (2) is true. When I entertain (1), considering a possibility in which it is true, could it be plausible that I am thereby considering worlds which are truly described by (2)? The idea that there even are worlds or situations truly described by (2) seems, frankly, confused. That is just the upshot of recognizing, as we did in the last section, that what it says is not, in any ordinary sense, supposable or conceivable.

The point can be made in reverse. Suppose (1) and (2) have the same truth-conditions. Then, since there is nothing that counts as coherently entertaining the possibility that (2) is true, there should be nothing that counts as coherently entertaining the possibility that (1) is true. Hence (1) is predicted to be unsupposable. That is the wrong prediction.

Unsurprisingly, the current theory also gets the relevant embedding facts wrong. According to this theory, these conditionals ought to have readings on which they are equivalent:

- (9) If it isn’t raining and I don’t know it, then there is a truth I don’t know.
- (27) # If it isn’t raining and it might be raining, then there is a truth I don’t know.

Manifestly they do not. (27) is not even interpretable.

So the account fails. Why bother arguing that an account already widely thought to be wrong is, in fact, wrong? The point here wasn’t to show that the speaker-relative account is wrong, so much as to demonstrate a particular way in which it goes wrong. For every other standard truth-conditional account of epistemic modals goes wrong, I think, in basically the same way. On any assignment of truth-conditions to ‘It might be

raining' that makes it a characterization of an epistemic state or states, the truth-conditions of 'It isn't raining and it might be raining' will be perfectly clear, and there will be no general difficulty with supposing or conceiving these conditions to obtain. The worlds in which it is true, in other words, will be clear to anyone with the relevant semantic competence. And, again, there will be no fundamental semantic obstacle to interpreting the corresponding clause in an indicative conditional, or under 'suppose'. These are the wrong predictions.

That is the basic objection. To further illustrate it, consider some ways of modifying the speaker-relative approach we have just considered. That approach might be modified along two parameters, corresponding to the following two questions:

- (Q1) Who is the epistemic agent (or set of agents) relevant to fixing the domain of quantification  $R$  for the modal?
- (Q2) What is the epistemic/evidential relation between that epistemic agent (or set of agents) and the set of propositions whose intersection fixes the domain of quantification  $R$  of the modal?

According to the speaker-relative account, the answer to the first question is *the speaker of the context of utterance  $c$* , and the answer to the second question is simply *knows* (more precisely, *knows at  $w$* ). A variety of alternative answers to these questions are available. For instance, in response to the first question, we might replace *the speaker of  $c$*  with

- the discourse participants of  $c$
- the epistemic agent(s) intended by the speaker of  $c$
- the "relevant community" at  $c$ <sup>12</sup>
- everyone in the world of  $c$

Etc. (Alternatively, instead of picking out the epistemic agent(s) in a way opaque to operators, one might also identify the epistemic agent(s) with the value of a shiftable index parameter, one that receives a default setting—perhaps one of the above—by the context.) As for the second question, we might replace *knows at  $w$*  with

<sup>12</sup>For this proposal see DeRose (DeRose, 1991, 593-4). It does not really answer the question so much as openly evade it. A second way to evade the question is to use the passive voice, speaking of "what is known" without being explicit about who is doing the knowing. This strategy occurs in Hacking (1967) and occasionally in Kratzer (1981). It seems to me that insofar as this question of who is doing the knowing goes without any systematic answer, one's account is to that extent incomplete.

- has good evidence for at  $w$
- knows, or has a relevant way of coming to know, at  $w$ <sup>13</sup>
- distributively knows at  $w$ <sup>14</sup>

and so on.

Now, it would be very tedious to apply the simple worry I am raising in detail to each possible setting of these two parameters. Perhaps it is already obvious that none of these options will generate the unpossibility we actually find in  $\ulcorner \neg\phi$  and MIGHT  $\phi \urcorner$ . But if not, here is a simple way to see why. Consider the following sentence:

(28) It is compatible with all our evidence that there is lead on Pluto.

where by “our” I intend (unrestrictedly) *everyone*, and where by “all our evidence” I mean all of it pooled together. So understood, (28) would entail, according to all permutations of the standard approach I am aware of, the sentence

(29) There might be lead on Pluto.

relative to a fixed context. Consequently, according to all these accounts, (30) transparently entails (31) relative to a context:

(30) There is no lead on Pluto and it is compatible with all our evidence that there is lead on Pluto.

(31) There is no lead on Pluto and there might be lead on Pluto.

But this is wrong. (30) may well be true, but not (31), for it doesn't make sense.

The only clear way I see of evading the worry I am raising while remaining within a standard truth-conditional approach is to make ‘might’ and ‘possibly’ factive. If we did that, then we would have a straightforward explanation for why sentences of the form  $\ulcorner \neg\phi$  and MIGHT  $\phi \urcorner$  do not express conceivable propositions: the reason is that they are just contradictions,

<sup>13</sup>Cf. DeRose (DeRose, 1991, 593-4).

<sup>14</sup>Cf. von Fintel and Gillies (2004). *Distributed knowledge*, a kind of group knowledge, is intended in the sense of Fagin, et. al. (Fagin et al., 1995, p.3): a group has distributed knowledge of  $\phi$  just in case “by pooling their knowledge together the members of the group can deduce  $\phi$ , even though it may be the case that no member of the group individually knows  $\phi$ .” The idea here is to get the domain of modal quantification simply by intersecting all the propositions known by each individual in the group.

false at every possibility. This move would certainly evade my objection, but, as already noted, it can be rejected for independent reasons. Possibly there is lead on Pluto, possibly there isn't; but from this we cannot infer that there both is and is not lead on Pluto. Few modals are better representatives of *non*-factivity than epistemic 'might' and 'possibly'.

## 5 Eavesdropping

I have argued that the proprietary weirdness of sentences like  $\ulcorner \neg\phi$  and MIGHT  $\phi \urcorner$  is a worry for standard accounts of epistemic modals. A different kind of objection to standard accounts has been urged recently by MacFarlane (2003), Egan et al. (2005), and Egan (2005). There is, I think, a simple point of connection worth mentioning between their worry and the one I have just offered.

Their worry concerns cases of *eavesdropping*.<sup>15</sup> Suppose that Sally is having a conversation with George, and that Jane is eavesdropping on them. Sally suspects that Joe is in Brooklyn; George has no idea where Joe is; and Jane, having just walked by Joe, knows that Joe is not in Brooklyn. Suppose now that the following things get said:

CASE 1

*Sally (to George)*: Joe might be in Brooklyn.

*Jane (to herself)*: That's not true; I saw Joe just a few minutes ago.

Here, as I understand it, the case is supposed to provoke two semantically problematic intuitions: (i) that the eavesdropper is making a judgment about the truth-value of Sally's sentence-token; and (ii) that the eavesdropper's judgment is correct (true). Insofar as we have those intuitions, the problem is supposed to be that standard approaches to epistemic modals can't plausibly accommodate them. For in order to accommodate them, it appears that the truth-conditions of Sally's claim now needs to be sensitive to the epistemic states of, not just members of some salient or intended community, but also potential eavesdroppers. That is, to get the truth-values right, we'd need to add potential eavesdroppers to the community relevant to the evaluation of Sally's claim. The trouble with doing this, as MacFarlane emphasizes, is that it would lead to truth-conditions for epistemic possibility claims that are implausibly strong, making such claims unassertable at contexts in which they are in fact assertable. We would,

<sup>15</sup>The case that follows is modified from MacFarlane (2003).

in effect, end up with a much-too-large relevant community, too large to think it plausible that Sally was ever really trying to hold forth on what is compatible with that group's knowledge.<sup>16 17</sup>

I think these eavesdropping cases are revealing and problematic for standard accounts, but I won't try to defend that view here. I just want to point to an apparent resemblance in strategy between the challenge posed by eavesdropping cases and the puzzle I have raised. In the eavesdropping cases, we stipulate a case wherein the proposition expressed by some sentence  $\phi$  is false, and then go on to ask, in various (direct and indirect) ways, about the truth-values of MIGHT $\phi$ -sentences as tokened in the

---

<sup>16</sup>It is worth emphasizing here that making the relevant epistemic community much larger than simply the speaker tends to yield totally implausible assertability conditions. So too does making the epistemic community "flexible" in ways out of the control or knowledge of the speaker: where would one get off making a claim concerning the knowledge of a group whose membership they were actually unaware of?

<sup>17</sup>An aside: what, incidentally, is minimally required to raise this puzzle? One might suspect that the core difficulty eavesdropping cases pose is generated by their having just the following two properties:

- (a) In the imagined situation, there is an epistemic agent who says, with warrant and in a natural context, MIGHT  $\phi$ .
- (b) In the imagined situation,  $\phi$  is false.

According to this view, eavesdropping *per se* is not actually a central part of the puzzle at all. We can, if we want, forget about Jane and just ask ourselves: *given* that Joe is not in Brooklyn, is Sally's sentence tokened true or not? In a way, of course, we already seem to be doing this indirectly, when we assess Jane's utterance. Insofar as we find it queer to answer this question "yes", we have a problem for the standard accounts. The situation may be further dramatized by adding to the mix a knowledgeable eavesdropper whose utterances we can have further judgments about, but (so goes this view) that further step is not essential.

It would be nice if this sort of view were correct—it would let us simplify the examples. But it is unclear that it could be the whole story. While most people do seem to find something about what Sally says here somehow wrong or off:

CASE 2:

(Stipulation: Joe is not in Brooklyn.)

*Sally (to George):* Joe might be in Brooklyn.

and while that judgment is itself *prima facie* problematic for standard accounts, it appears that introducing a knowledgeable eavesdropper into the picture (Case 1) just seems to generate a stronger judgment to the effect that what Sally said is false. We should wonder why exactly this is. After all, you already know (in Case 2) what Jane knows (in Case 1). You know it, in fact, with the awesome power of stipulation. We might then expect your judgment to be all the stronger.

imagined case. In the puzzle I describe, we stipulate that the proposition expressed by  $\phi$  is false, and attempt to further stipulate that MIGHT  $\phi$  is false. Both puzzles seem to result, at least in part, from our reluctance to regard both  $\neg\phi$  and MIGHT  $\phi$  as true of any scenario.

## 6 Assessment-sensitive semantics

I am not sure, however, that those theorists who have emphasized eavesdropping would tend to see the situation in quite this way. I am not sure because their positive proposals for epistemic modals—versions of what they call a *relativist* or *assessment-sensitive semantics*—are prone to basically the same objection I have pressed against standard accounts. To spell this out, let me sketch a version of the relativist semantics, one that I think captures the core idea behind these proposals.

Reflecting on the way that eavesdroppers seem to evaluate the truth of overheard epistemic modal statements, a natural thought to have is that we tend to assess assertions containing epistemic modals “against our current epistemic situations, rather than the epistemic situations in which they were made” (MacFarlane, 2003, p.3). Supposing that thought is on the right track, we should then want to ask: how can a semantics for epistemic modals possibly give sense to this tendency? The relativist answer is that the way to make sense out of this tendency is to design a semantics that will enable the truth-value of a sentence token to vary systematically with the epistemic state of the agent assessing the claim for truth.

Here’s how that is done. First, recall the (closest thing to the) standard picture of what fixes the truth of a sentence token at a context of use:

A sentence  $\phi$  in a context of use  $c$  is true just in case the content expressed by  $\phi$  in  $c$  is true when evaluated with respect to the index determined by  $c$ . (paraphrasing Kaplan (Kaplan, 1989, p,522))

Or, to draw it out (reading the arrows as “determines”):

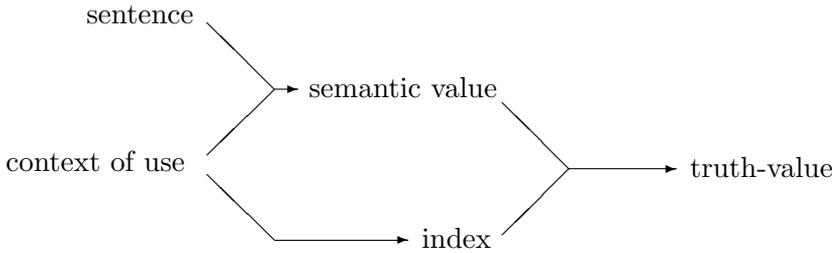


Figure 1: the “classic” picture of the relation between sentence, context, and truth.<sup>18</sup>

This is the picture built into the classic work of Kaplan (1989), Stalnaker (1970), and Lewis (1980). The context of use comes into play twice over on the road to truth here: first in determining content or semantic value, second in determining the index with respect to which that content is evaluated.

The relativists want to change this picture. Specifically, they want to drop the assumption that the context of use alone determines the index of evaluation. They want to allow instead for cases in which the index relevant to evaluating the truth of a sentence tokened at a context of use is *jointly* determined by both the context of use *and the context from which the sentence token is being assessed for truth*. Following MacFarlane, we can call the latter kind of context a *context of assessment*. Here then is the new picture, diagrammatically:

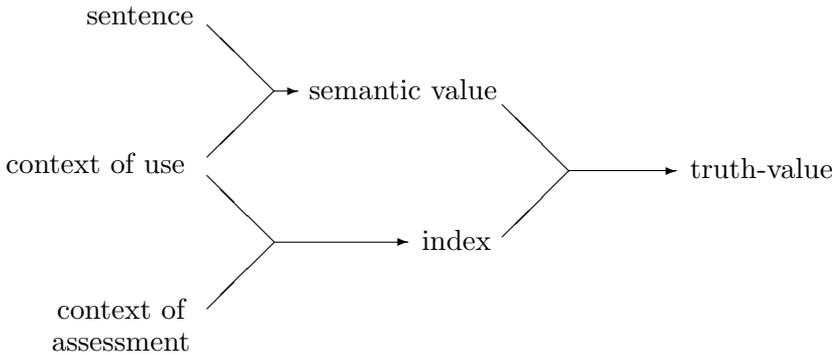


Figure 2: the relativist picture of the relation between sentence, context, and truth.

<sup>18</sup>Cf. Lewis (1980). I have added an extra arrow to Lewis’s original picture, the one indicating that it is the context of use which supplies the index of evaluation.

This picture enables sentence tokens to be (indexically) *assessment-sensitive*: their truth-values can in principle vary according to features of the contexts from which they are assessed. If there are sentences which do so vary, it immediately follows that the generalization:

Sentence tokens have their truth-values absolutely.

is false, assessment-sensitive sentences being straight counterexamples.

Notice that the *ways* in which a sentence-token can be assessment-sensitive depend entirely upon the parameters of the index assumed. If we restrict the index so that it contains only a world parameter, the relativist proposal is perhaps not all that exciting: the above absoluteness generalization may fail, but merely because the truth of some sentence token depends on the world at which it is assessed. The view gets interesting, and palpably relativistic, only when the index is enriched to include features which vary *within* worlds, so that we have intra-world variation in truth-value for a single sentence token.

It is intra-world truth variation that relativists about epistemic modals are after, and which they want to put to work explaining the eavesdropping cases. This kind of truth variation can be achieved in two steps. First step: endorse figure 2. Second step: enrich the index so as to include, in addition to a world parameter  $w$ , a time parameter  $t$  and an agent (or place) parameter  $x$ —that is, let the indices of evaluation be centered worlds. Then we can define a relativist proposal for epistemic ‘might’ as follows (letting ‘ $u$ ’ designate context of use, ‘ $a$ ’ designate context of assessment, and letting ‘ $\mathbf{c}$ ’ range over centered worlds  $\langle w, t, x \rangle$ ):

$$\llbracket \text{MIGHT } \phi \rrbracket^{u,a} \text{ is true at } \mathbf{c} \text{ iff } \exists \mathbf{c}' \in R_{u,a}(\mathbf{c}) : \llbracket \phi \rrbracket^{u,a} \text{ is true at } \mathbf{c}'$$

where the value of the set of centered worlds quantified over is defined as

$$R_{u,a}(\mathbf{c}) = \text{set of centered worlds compatible with what the assessor at } a \text{ knows at } \mathbf{c}$$

The proposal aims to accommodate the eavesdropping cases by allowing eavesdroppers to assign truth-values to overheard sentences tokens in a way that turns directly on *their own* epistemic state at the time. When Sally says above ‘Joe might be in Canberra’, both her sentence token and the content it expresses are true with respect to her context of assessment, and both are false with respect to Jane’s context of assessment. The very same content, and the very same sentence token, are true at one location in the world and false at another location in the very same world.

Notice that the definition of the restriction here speaks of the *centered-worlds* compatible with what is known, rather than simply of the worlds compatible with what is known. That is because the relativist proposal goes naturally with the view, defended by Lewis (1979), that the objects of the attitudes are *properties* representable by sets of centered worlds. Indeed some relativists do explicitly favor a relativist semantics + self-locating attitudes package deal.<sup>19</sup> I will assume this package, but be advised that the semantics above can accommodate a more traditional, worlds-only view. (For this, just add the stipulation that a centered world *c* is compatible with what *x* knows iff the world *w* of *c* is compatible with what *x* knows.)

The basic problem I posed for standard accounts recurs for the assessment-sensitive proposal just described. On the self-locating view of the attitudes, to stand in the supposition relation to a property is just to suppose oneself to be located at a center falling under the property. Thus, given the above assessment-sensitive semantics, to suppose that (2) is true (repeated):

(2) It isn't raining and it might be raining.

is to suppose oneself to have the property of *being located at a centered world where it isn't raining and not knowing that it isn't raining*. But that is a perfectly coherent, supposable, conceivable possibility—unlike (2), alas. So we have here the usual objection. In a way, the prediction of the relativist semantics is quite similar to that of the speaker-relative analysis within the standard approach, for in both cases supposing what (2) says is predicted to be basically the same thing as supposing what (1) says:

(1) It isn't raining and I don't know it isn't raining.

which we know cannot be right. (The two proposals would be still more similar, of course, if we dropped the self-locating attitudes picture in favor of worlds-only content.)

The objection I have raised is to a specific implementation of the assessment-sensitive approach. It remains to be seen whether a version of this approach different from the one I have described can overcome the objection.

## 7 Changing direction

The situation is as follows. Our troublemaking sentence 'It isn't raining and it might be' is not a contradiction, but neither does it describe a coherent

<sup>19</sup>See Egan et al. (2005) and Egan (2005). (I should emphasize that MacFarlane (2003) does *not* exploit centered worlds. I am not sure what he takes the relation between the semantic values of sentences and the objects of belief to be.)

possibility in the way Moore-paradoxical sentences do. It is hard to see how these two properties can both be accommodated within an approach assigning ‘It might be raining’ truth-conditions. For if it has truth-conditions, then either these conditions are compatible with the truth of ‘It’s raining’, or they are not. If they are compatible, then it is puzzling why the conjunction cannot be coherently supposed (conceived of, considered...). If they are not compatible, then MIGHT  $\phi$  entails  $\phi$ —a totally implausible result. We have some motivation, then, for reconsidering the assumption that epistemic modal statements have truth-conditions; or, more specifically, motivation for reconsidering the assumption that epistemic modals designate functions into propositions. But if epistemic modals don’t contribute to the determination of a truth-conditional content, what do they do?

Here is an old idea: what epistemic modals do is modulate *force*, not content.<sup>20</sup> POSSIBLY  $\phi$ , MIGHT  $\phi$ , PROBABLY  $\phi$ , LIKELY  $\phi$ —these all provide various ways of “putting forward” the proposition expressed by  $\phi$  with a force in some sense “weaker” than that of assertion. Epistemic modals do not, on this view, serve to characterize various modes in which a proposition is true with respect to a set of worlds. Rather, they provide “modes of putting forward” propositions as true, modes with varying associated degrees of strength (weaker for ‘might’, stronger for ‘probably’ and ‘must’).<sup>21</sup>

The idea is obviously very sketchy as described. But it is not unintuitive, and I think one can already see in it the vague outline of a solution to our problem. Conjunctions like ‘It isn’t raining and it might be’ would be viewed as asserting a proposition while simultaneously putting forward, to a weaker degree, its negation, the net yield being incoherence. The puzzles

---

<sup>20</sup>Old because it seems to have been fairly common in logic textbooks going back, at least, to the late 19th and early 20th centuries. See, for instance, the remarks in Jevons (Jevons, 1878, p.70) and Kneale (Kneale, 1949, p.3). A somewhat more recent advocate of this sort of idea, for ‘probably’, was S. E. Toulmin, who claimed that “...to say ‘Probably  $p$ ’ is to assert, guardedly and/or with reservations, that  $p$ ...” (Toulmin, 1956, p.190). Toulmin may have been following Wittgenstein, who, in discussing the meaning of ‘It may be raining’, gnominically warned: “Don’t regard a hesitant assertion as an assertion of hesitancy” (Wittgenstein, 1953, p.192).

<sup>21</sup>Here is another way of putting the thought. It is a commonplace that *to say what you believe* is not *to say that you believe it*. Rather, you just express the proposition you believe, and the information that you believe that proposition is carried by the assertive force with which you express it. A force approach to epistemic modals tries to extend this commonplace to partial belief. To say what you *partially* believe—with an epistemically modalized claim—is not to say that you partially believe it. Rather, you just express the proposition you partially believe, and the information that you partially believe that proposition—to whatever extent you partially believe it—is carried by the (yet to be explicated) non-assertive force with which you express it.

about eavesdropping also seem to get the beginnings of an explanation: if saying POSSIBLY  $\phi$  or PROBABLY  $\phi$  are varying ways of putting forward the proposition expressed by  $\phi$ , then what better informed eavesdroppers are doing, when they overhear these remarks and judge them false or mistaken, is simply denying the proposition being put forward, viz. the one expressed by  $\phi$ .

This looks at least mildly promising. It is worth mentioning, in addition, that a force account of epistemic modals would seem to fit rather naturally with the way that Bayesians tend to describe epistemic modal beliefs. The usual way of modeling, within a Bayesian framework, someone who we would describe as believing (e.g.) VERY PROBABLY  $\phi$  would be to let the credence function characterizing their credal state map  $\phi$ , or the proposition it expresses, to some high value. Whether someone accepts what an epistemic modal statement says is thus generally taken to be a matter of their credence in the proposition expressed by the sentence embedded under the modal. So whether ‘Jones believes that it’s very probably raining’ is true is a matter of Jones’s credence in the proposition that it’s raining. In the attitude report, the modal tends to be treated, as it were, adverbially: the object of Jones’s attitude is the proposition that it’s raining, and the modal tells how strongly the proposition is believed.<sup>22</sup> (The object of the attitude is *not* taken to be some second-order belief concerning what is true about some epistemic states, as would be expected if one of the standard truth-conditional accounts discussed above were correct.) A force account of epistemic modals would corroborate this Bayesian tendency: in the unembedded case, it treats the epistemic modals as modulators of “assertive force”, just as the Bayesians tend to treat them, in attitude reports, as modulators of degree of belief. Insofar as this Bayesian tendency strikes us as reasonable, a force account gains credence.

Of course, vague outlines of solutions are not solutions. The preceding has just by way of motivating an attempt to fill in the outlines. It remains to spell out a force analysis of epistemic modals in detail, and how to do this is not at all obvious. Indeed, one might worry about a view that requires so many scare-quotes to describe. We need a story about what “putting forward” could possibly be, how it can come in degrees, in what sense these kinds of force might be comparatively weaker than assertion, in what sense force can even be the sort of thing that attaches to a lexical

<sup>22</sup>This descriptive tendency is not difficult to find even in early work on the notion of credence. For instance: “In forming any judgment we cannot avoid attaching to it a particular degree of credence, which might be, and often is, expressed by the insertion of some adverb to qualify the copula; thus ‘Tomorrow will (possibly) be fine’...” (Thomson, 1877, p.278).

item, how such lexical items can embed, and so on. Maybe, in the end, we will decide that the complexity of a force account is not worth the price. Maybe—but we cannot judge the cost until a reasonably detailed account is put on the table. What follows is a first attempt at constructing such an account, at making a force account of epistemic modals precise. It is highly preliminary, and only intended as a beginning.

## 8 Assertion revisited

The objective is to see epistemic modals as modulators of assertive force. We therefore need to see assertion as the sort of thing that either comes in degrees, or is itself located on some scale. To get where we want to go, it will help to take as our starting point a certain standard account of the force of assertion. Then we can ask, of that account, whether it will allow us to see assertion in this light, and if not, what modifications are required.

According to the influential view I will start with, the “essential effect” of assertion is “to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed” (Stalnaker, 1978, p.86). On this picture, what is presupposed in a discourse context is represented by a *context set*, the set of possible worlds compatible with the shared presuppositions of the discourse participants. The context set can be seen as a theoretical representation of the context itself as the discourse participants understand it. Assertions are either accepted (successful) or rejected (unsuccessful). The information growth achieved by a successful assertion is represented by intersecting the propositional content of the assertion—a set of worlds—with the context set. If an assertion is rejected, the context remains unchanged (*modulo* the addition of the information that the relevant assertion was made and rejected). Each discourse participant has her own context set—her own representation of the context—but in the non-defective case, these are all identical, and we can speak of *the* context set.

Think of it like this. Let  $\mathbb{C}$  be the domain of context sets. Whenever it is recognized that a proposition  $p$  is being asserted, the application of a certain function—call it **assert**—is triggered<sup>23</sup>:

$$\text{assert}(p) = f : \mathbb{C} \rightarrow \mathbb{C}, \text{ where for all } c, f(c) = c \cap p$$

The function **assert** maps propositions into context-change potentials, functions from contexts into contexts. Alternatively, we can uncurry the func-

<sup>23</sup>Assuming, that is, the assertion is not rejected.

tion and see it as mapping  $\langle \text{proposition, context set} \rangle$  pairs into context sets:

$$\text{assert}(p, c) = c \cap p$$

Call a function with the type of `assert` a *change function*. When the speaker makes an assertion, then, she is “proposing” that you take the proposition she has expressed together with the context set, apply `assert`, and let the result be the new context set. Consider `assert` a formal representation of assertive force. It is a sort of instruction telling you what to do with the proposition expressed.

To diagram things:

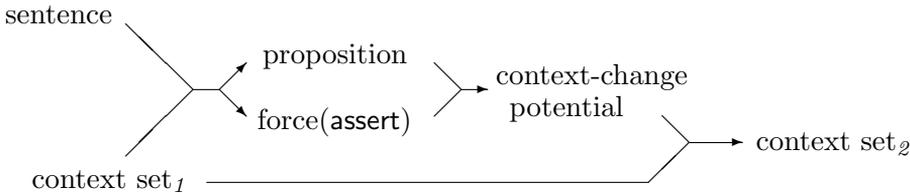


Figure 3: Assertion on Stalnaker’s model.

Sentence and context indicate, not only the proposition expressed, but also `assert`. These yield a context-change potential which, applied to the previous context set, yields a new context set updated with the proposition expressed. (Note that context set<sub>1</sub> includes the information that the relevant sentence has been uttered.)

I have characterized Stalnaker’s model in a somewhat non-standard way in order to make salient a candidate for the meanings of the epistemic modals. The idea is to treat them as literally designating functions on par with `assert`—functions with the same domain and range, but which yield different context-change potentials. If plausible change functions for these modals could be found—ones yielding context changes that are sort of like assertion, only “weaker”—this would give us a fairly straightforward way of seeing epistemic modals as modulators of the force with which the propositions they embed are expressed.

Of course, as we’ve seen, epistemic modals can be embedded. So if we want to take this kind of approach, a picture like Fig. 3 won’t exactly do as it stands, for this picture represents the change function as always applying last. The picture needs to be a more general one:

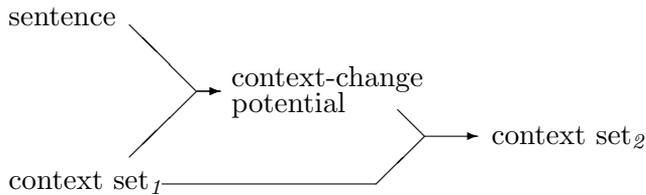


Figure 4

where we see the context-change potential of a sentence in context as compositionally determined, either via the application of a change function to a proposition, or as a function of several context-change potentials each of which is itself so determined.

I will pursue an picture along these lines in what follows. Before I start considering specific candidate change functions for the epistemic modals, however, I need to make two clarificatory remarks about assertion to set the stage.

First, there is a basic question about Stalnaker's model that needs answering, but whose answer I will not give here. The question is: what feature of sentence or context is it, exactly, that indicates *assert*? We have a reasonably worked-out program for saying how the relevant proposition is indicated: it's called truth-conditional semantics. But how is known that the proposition expressed is being asserted? Exactly what about the sentence or the context carries this information? The question does not seem to me to be widely recognized, and among those who recognize this question, there seems to be no consensus concerning its answer.<sup>24</sup> The question matters here for an obvious reason. We are pursuing the idea that epistemic modals designate functions like *assert*. Since epistemic modals are lexical items in the syntax, one might expect *assert* to correspond to some element in the syntax as well. But if so, what is this element? And if it is not designated by some element in the syntax, what represents it? I will not directly broach these questions here. They are important, but in a sense they are premature. I will just put *assert* where it appears to be theoretically needed. Once it becomes reasonably clear what the distribution of *assert* needs to be, we will be better placed ask whether there is any syntactic object or property matching this distribution.

<sup>24</sup>Grammatical mood seems to be a frontrunner as a representer of force—but then, there seems to be little agreement about precisely what mood is.

Second, I want to emphasize that I am not trying to perform conceptual analysis on the concept of assertion, or on the concept of force. Nor do I assume a notion of force according to which it is fundamentally a component of action under intentional description. I do not assume that assertion is a speech act on par with a rational action like making a promise or a warning. To assert is just to revise the context set in a certain way; and it may be that revising the context set, unlike promising, is not really the sort of thing you do intentionally. I wish to leave the issue open. It may be that revising the context set is no more intentional than performing semantic composition during interpretation, or stopping the airflow in your vocal tract in order to generate the bilabial plosives in “pop”. No doubt you intentionally convey information, but do you intentionally update the context set? Perhaps; but I will not assume as much. As the context changes we define get more complicated, and as the change functions we consider get more deeply embedded, we may find it hard to see the specific context changes invoked as intentional in any rich sense. In keeping with my agnosticism on this matter, I will avoid the metaphor of assertion being a kind of *proposal* which an interlocutor can, as a matter of decision, *accept* or *reject*. It will be simpler to just see assertion as having its “essential effect” on the context set automatically. What we would intuitively call the “rejection” of an assertion will be seen, not as blocking the effect of assertion, but rather as (automatically) reversing this effect.

## 9 Initial attempts at a force analysis

Returning now to the main thread, the thought is that the epistemic modals designate change functions like **assert**, the difference being that the context changes these modals generate are, in some sense we now need to formally spell out, weaker. So how should these functions be defined?

One thought is that the functions have something like this form:

$$\text{epistemic modal}(p, c) = s \cap p$$

(where  $s$  is some subset of  $c$  determined by the modal)

Epistemic modal statements would have an effect like assertion, but the propositions they embed would intersect with only a subset of the context set, rather than the whole context set. The subset would be smaller for the “weaker” epistemic modals (‘might’, ‘possibly’, etc.), thereby eliminating fewer worlds, and larger for the “stronger” epistemic modals (‘probably’, ‘very likely’ etc.), thereby eliminating more worlds. Intuitively, it would be as if epistemic modal statements tacitly enforce a bunch of presuppositions

(thereby shrinking the context set), make an assertion with respect to this restricted context set, and then release the presuppositions, returning to the context the worlds temporarily eliminated. (They would have an effect similar to that of indicative conditionals, at least on one very context-sensitive analysis of them.) In short, epistemic modals would modulate the force of assertion by temporary restricting the context set to various degrees.

The only worry I have about this account is that I don't know how to finish it. It is hard to see how to define the needed restriction for any given epistemic modal. What is  $s$ —what is it that gets tacitly and temporarily presupposed? Perhaps what we need to say that what gets tacitly and temporarily presupposed are propositions that span the evidential gap between the speaker's knowledge and the proposition embedded under the modal. That is, perhaps the temporarily presupposed propositions are those propositions which, if the speaker knew them, would put him in a position to assert  $\phi$  without qualification. But normally, for any given epistemic modal, a huge variety of propositions will span this gap, making the restriction radically underdetermined. Because I see no natural way to resolve this issue, I will not consider this proposal further.

A second idea in the vicinity of a force account, due to Veltman (1996), concerns epistemic 'might'. According to Veltman, MIGHT  $\phi$  statements do not change contexts in anything like the way that assertions do. Rather, they "provide an invitation to perform a test" on the context set—a test, roughly, for whether the context set contains  $\phi$ -possibilities. Adjusting the technical implementation for our setting, the proposal is that 'might' designates the function *might*, which behaves as follows:

$$\text{might}(p, c) = \begin{cases} c & \text{if } c \cap p \neq \emptyset; \\ \emptyset & \text{otherwise.} \end{cases}$$

Assuming there are possibilities compatible with the proposition  $p$  expressed by  $\phi$  in the context set, the context-change potential associated with MIGHT  $\phi$  makes no change to the context set. (If there are no such possibilities, then—as when an assertion is made that is incompatible with the shared presuppositions—the empty set results, presumably calling for some kind of repair strategy.) This certainly looks like a much weaker effect than that of assertion. The idea, I take it, is that the point of saying MIGHT  $\phi$  is to ensure that  $\phi$ -possibilities are treated as compatible with the presuppositions of the context.

I actually think this proposal can, with some additional assumptions, explain a number of the puzzling features of sentences like  $\ulcorner \neg\phi$  and MIGHT

$\phi^{\neg}$ . The basic idea, roughly, would be that the two conjuncts require mutually incompatible context changes, so that there is no single context in which both conjuncts are accepted. I also think this account can give a reasonable explanation of the eavesdropping cases. The basic idea, roughly, would be that the eavesdropper who knows  $\neg\phi$  judges an overheard MIGHT  $\phi$ -sentence false because the “test” corresponding to this sentence fails (i.e., yields the empty set) at his context, given that he is disposed to presuppose  $\neg\phi$ . I will not stop to elaborate or defend these positions here, however, nor will I advocate this proposal. I mention Veltman’s approach only in order to point out that it has a basic limitation: there is no clear way to extend it gradable epistemic modals such as ‘possibly’, ‘probably’, and ‘it is likely that’. Consider:

(32) It is very, very possible that I left the stove on.

(33) It’s probably snowing in Cambridge.

(34) It is extremely unlikely that there is lead on Pluto.

It is hard to see what the modifier ‘very’ could possibly be doing in (32), if the ‘possibly’ operator, like ‘might’, is just there to invoke a test. It is correspondingly hard to see how to extend Veltman’s proposal to epistemic modals such as ‘probably’ and ‘likely’. What change functions correspond to these modals? What tests, if any, correspond to them, and how are they changed by modifiers? Here it becomes clear that Veltman’s proposal does not quite fit into the epistemic-modals-as-assertive-force-modulators picture I have been trying to push, for the proposal does not give us a parameter to modulate. Without that parameter, I do not know how to triangulate my way from might and assert to possibly, probably, and the rest. Since I want to explain what is puzzling, not only about

(2) It isn’t raining and it might be raining.

but also about sentences like

(35) It isn’t raining and it is very possibly raining.

(18) It isn’t raining and it is probably raining.

I will move on.

A more promising strategy for articulating a notion of “degrees of force”, I think, would be to take a cue from the Bayesians. Recall that the Bayesian characterization of an arbitrary belief will consist of at least

two components. There is, first, the content of the belief, or what the belief represents. This is modelled in the usual way by a proposition, a set of worlds. Second, there is the strength or degree to which the proposition is believed. This is modelled by assigning the proposition a real number between 0 and 1, a number determined by the agent's *credence function*, a probability measure defined over possible worlds representing the agent's (idealized) credal state. Thus on the Bayesian picture a given credal attitude involves two sorts of distinction. First we need a proposition, a way of distinguishing between possibilities; and second we need a credence function, a way of distinguishing between (subjective) probabilities of possibilities. Again: one thing to distinguish between possibilities, another thing to distinguish between the probabilities of possibilities. Bayesians wheel in credence functions to do the latter. Now suppose this picture is not wholly confused, and suppose further that epistemic modal statements serve, in the first instance, to distinguish the probabilities or likelihoods of possibilities, and not merely to distinguish between possibilities, as advocates of standard truth-conditional approaches maintain. Then it would seem that, in order to model the effects of these statements—the changes they make to communicative contexts—we need an additional layer of information in our representation of the context, one playing something like the role that credence functions play for the Bayesians.

It seems to me that the basic assumption required by such an approach would be, mirroring the Bayesian approach to belief, something like this:

ACCEPTANCE BY DEGREES: propositions may be accepted in a discourse context to various degrees.

This is as opposed to the usual notion of acceptance, where propositions are viewed as either accepted or not accepted in context (full stop). I will spend the balance of this paper trying to work this vague thesis of acceptance by degrees into some more specific form. The direction I take I will call *context probabilism*:

CONTEXT PROBABILISM: a discourse context should be modelled (at least in part) by a probability measure over possibilities, or by a set of such measures.

Like the thesis of acceptance by degrees, context probabilism might be developed in a number of ways. The path I take is surely not the only possible one.

## 10 Context probabilism I

An obvious first thought at pulling off context probabilism would be to try defining a single probability measure over a space of worlds, letting that measure serve as the fundamental representation of the discourse context. Let's start with that.

Assume that every discourse participant has a representation of what I will call the *context probability space*. Let this play the role earlier played by the context set. When things are going well, the representation of this space is the same for all discourse participants. A context probability space is a probability space  $\langle c, P, Cr \rangle$  consisting of

1. A set of worlds  $c$  (our “sample space”)
2. A  $\sigma$ -field  $P$  over  $c$
3. A *context credence function*  $Cr : P \rightarrow [0, 1]$ , where  $Cr(c) = 1$

The set of worlds  $c$  is the space of possibilities that are candidates for being distinguished between in the conversation. The elements of  $P$  are propositions. Assuming for simplicity that  $c$  is countable, we can just think of  $P$  as the power set of  $c$ —as the set of all possible ways of dividing the possibilities in  $c$ . The *context credence function*  $Cr$  characterizes the degree to which each proposition in  $P$  is accepted in the discourse context. It maps each element of  $p$  to a real value representing what I will call the *context credence* in  $p$  (for short, “c-credence”). Just as credence is stuff of (partial) belief, context credence is the stuff of (partial) acceptance. The higher the real value of  $Cr(p)$ , the more  $p$  is accepted. In this framework, it is fundamentally a c-credence function, not a set of worlds, that represents the context. A c-credence function is, formally, a probability measure. The motive for developing a notion of acceptance by degrees by using probability-theoretic tools—that is, via context probabilism—is simple: it is to preserve certain intuitive logico-probabilistic relations between propositions in context, relations it seems natural to assume are in force. The assumption ensures, for instance, that the c-credence of  $p$  in context goes up as the c-credence in  $\neg p$  goes down.

I distinguish c-credence from credence because giving a model for communication is one thing and giving a model for belief update and revision is another. Acceptance—the propositional attitude I am saying we take towards the body of information in play during communication—is similar to belief in many respects, but the two are not the same. One can accept things in conversation that one doesn't actually believe, and vice versa.

Our working account of the effect of assertion—that is, Stalnaker’s—was not designed to interact with this kind of representation of context. Perforce we need to modify it. How? Earlier, where our representation of the context was taken to just be a set of worlds, the essential effect of assertion was to eliminate possibilities from that set. We can give an analogue of that effect here, by modeling assertion as *conditionalization*. Conditionalization is a method for transforming (“updating”) probability measures. The essential effect of assertion, in this model, will be to update the c-credence function by conditionalizing on the asserted proposition. If  $Cr_{\text{OLD}}$  is the c-credence function of the context before the assertion, and  $p$  is the proposition just asserted, then for any proposition  $q$ , the new c-credence in  $q$  will be given by  $Cr_{\text{NEW}}$ :

$$Cr_{\text{NEW}}(q) = Cr_{\text{OLD}}(q|p)$$

(Where  $Cr_{\text{OLD}}(q|p)$ , the previous conditional c-credence of  $q$  given  $p$ , is defined in the usual way via the ratio formula as  $Cr_{\text{OLD}}(q \wedge p) / Cr_{\text{OLD}}(p)$ , where  $Cr_{\text{OLD}}(p) > 0$ .) Our assert function is now a function from propositions and c-credence functions into c-credence functions:

$$\text{assert}(p, Cr_{\text{OLD}}) = Cr_{\text{NEW}}$$

An asserted proposition become, as we might put it, *taken for granted*: the c-credence in that proposition goes to 1. Note that if  $p$  is asserted but  $Cr_{\text{OLD}}(p) = 0$ ,  $Cr_{\text{NEW}}$  will be undefined. This is a way of representing the fact that  $p$  is not compatible with contexts where  $\neg p$  is taken for granted.

A metaphor popular in Bayesian pedagogy may be helpful here. Start with a representation of logical space. (Recall that logical space is shaped like a rectangular box.) Now imagine smearing some fixed amount of mud over that space. The mud represents credence. The more mud in a region, or over a world, the more the agent represented is said to believe that the actual world lies within that region, or is that world. When you draw ovals on this logical space, representing propositions, the credence in that proposition is measured by the sum total of mud falling within its boundaries.

Now transplant this metaphor over to the current model of assertion. Our space is  $c$ . The space has a fixed amount of mud distributed over it, representing c-credence. When  $p$  is asserted, all the mud on the  $\neg p$ -worlds is moved, proportionally, to the  $p$ -worlds. As a result, every world with any amount of c-credence in  $c$  is a  $p$ -world; the  $\neg p$ -worlds lose all their c-credence. Assertion eliminates possibilities on this model—it just does so by pushing context credence off of possibilities and onto others,

in accordance with conditionalization. Assertion is a matter of pushing c-credal mud around.

Now to the epistemic modals. Again, our modals are to be functions like *assert*. In this setting, it is fairly easy to see how epistemic modals might be modelled as assertive force modulators: whereas asserting  $p$  would raise the context credence in  $p$  to degree 1, putting that proposition forward to weaker degrees, via epistemic modals, would assign  $p$  some weaker value, depending on the strength of the modal. In order to characterize the context change induced by these modals, we'd need a more general principle than conditionalization, for conditionalization only tells us how to revise the c-credence function in the special case that a proposition gets assigned degree 1 c-credence. The generalization I will use is *Jeffrey conditionalization*. If you Jeffrey conditionalize on  $p$ , where  $p$ 's new c-credence is  $Cr_{\text{NEW}}(p)$ , then your new  $Cr$  for any  $q$  is defined as

$$Cr_{\text{NEW}}(q) = Cr_{\text{NEW}}(p) \cdot Cr_{\text{OLD}}(q|p) + Cr_{\text{NEW}}(\neg p) \cdot Cr_{\text{OLD}}(q|\neg p)$$

The idea is then be that the epistemic modals would assign the proposition  $p$  they embed some real value (generally  $\neq 1$ ) and invoke Jeffrey conditionalization on the context with that value. To give a toy example, suppose we say that *probably* (the function expressed by 'probably') puts the c-credence of  $p$  at a context to 0.7. Then the value of  $Cr_{\text{NEW}}(p)$  would be 0.7, and *probably* would trigger Jeffrey conditionalization as follows:

$$\text{probably}(p, Cr_{\text{OLD}}) = Cr_{\text{NEW}}$$

with  $Cr_{\text{NEW}}$  defined as

$$Cr_{\text{NEW}} = f : P \rightarrow [0, 1] \text{ such that} \\ \text{for any } q \in P, Cr_{\text{NEW}}(q) = (0.7) \cdot (Cr_{\text{OLD}}(q|p)) + (0.3) \cdot (Cr_{\text{OLD}}(q|\neg p))$$

The other epistemic modals would similarly invoke Jeffrey conditionalization, depending on the new real value they assign to the propositions they embed.

While this approach is still rough and sketchy as described, it already seems to me to have quite non-trivial defects. Fortunately, the defects point the way towards a more plausible version of context probabilism.

Here are the defects. I called the case with 'probably' just described a toy example because it is not very plausible that *probably* puts the c-credence of  $p$  in a context to 0.7. The reason it isn't plausible is just that 0.7 seems to be totally arbitrary. The first worry is that this kind of arbitrariness is unavoidable on the current approach, for it requires that we

correlate each epistemic modal with some real number between 0 and 1. That won't do. A context in which  $p$  is accepted to degree 0.7 sounds like a context where PROBABLY  $\phi$  is accepted, but many other contexts are just as plausible candidates. What about contexts where  $p$  is accepted to degree 0.6? Or to 0.9145? Or to 0.55? It would be more sensible to say that PROBABLY  $\phi$  is already accepted in *all* these contexts, since in all of them  $p$  is accepted to a greater degree than  $\neg p$ . The problems multiply when we attempt to assign  $Cr$  values to modals like 'might' and 'unlikely'. What value goes with 'might'? A low value might be the first guess; surely a high value is not required. But shouldn't we want to say that, if PROBABLY  $\phi$  has been said and is accepted in a context, then *ipso facto* what MIGHT  $\phi$  says is also accepted in that context? We should; but there is no way of making sense of that thought on the current approach. If a given context is in the range of  $\text{probably}(p, Cr)$ , then it will just fail to be in the range of  $\text{might}(p, Cr)$ , assuming these modals assign  $p$  differing credences. Relatedly, what about 'unlikely'? Presumably that would get a low value, like 'might'. But 'might' and 'unlikely' obviously differ significantly in meaning; the latter makes for an intuitively stronger claim. How can this difference be captured just by assigning these modals different real numbers?

That makes two related worries. First, the framework seems to require totally arbitrary decisions about which real values to correlate with which epistemic modals. Second, the framework fails to capture the logical relations between these modals. If EPISTEMIC MODAL<sub>1</sub>  $\phi$  intuitively entails EPISTEMIC MODAL<sub>2</sub>  $\phi$ , we want it to be the case that the contexts updated by EPISTEMIC MODAL<sub>1</sub>  $\phi$  already reflect the information change which would be induced by EPISTEMIC MODAL<sub>2</sub>  $\phi$ . We don't yet have that here.

To overcome these weaknesses, I propose the following enrichment: let the discourse context be modelled as, not a  $c$ -credence function, but as a *set* of  $c$ -credence functions.<sup>25</sup> Propositions will thereby be associated with, not a single value, but a range of values, one for each  $c$ -credence function in the set of such functions representing the context. We can then treat assertion and epistemic modals statements alike as varying ways of constraining the range of values a proposition gets assigned in context. As I will set things up, they will do this fundamentally by *eliminating*  $c$ -credence functions from the set representing the context. This extension will let us evade the arbitrary real number assignments to the modals, and as we will see, it will give us the entailment relations we want via ordinary set-theoretic relations.

<sup>25</sup>This enrichment has its analogue in Bayesian models of belief. See for instance Jeffrey (1987), Kaplan (1996), and the numerous references cited therein.

One might expect this sort of enrichment of the context to make the implementation of context probabilism an order of magnitude more complicated. On the contrary, however: things are simpler on this approach. Let me lay out the details.

## 11 Context probabilism II

### BASICS

A discourse context is characterized by a set  $C$  of context credence functions  $Cr$ , which we call a *probabilistic context set*. A probabilistic context set represents both a set of worlds—the worlds which are in the domain of the  $Cr$  functions it contains<sup>26</sup>—and a collection of c-credence distributions over these worlds, or over the corresponding  $\sigma$ -field. If a given c-credence function is in a given probabilistic context set, call it *admissible* with respect to that set.

Suppose that the truth or falsity of some proposition  $p$  is logically independent of whatever has been accepted in a given conversation. Then, for every real number  $r$  such that  $0 \leq r \leq 1$ , the probabilistic context set  $C$  representing this discourse will contain a  $Cr$  such that  $Cr(p) = r$ . This represents the fact that the context in question is just *indifferent* about  $p$ : since nothing that has been said places any constraint on the possible values of  $Cr(p)$ , no value is excluded. The discourse participants don't yet jointly accept anything incompatible with any particular assignment of context credence to  $p$ .

Now suppose that  $p$  is asserted. As earlier, I will take it that assertions just become taken for granted in conversation. In the current setting, I will take this to mean that an assertion of  $p$  imposes a requirement on the members of the probabilistic context set: the requirement is that these functions map  $p$  to 1. Hence I define the *assert* function as follows:

$$\text{assert}(p, C) = \{Cr \in C : Cr(p) = 1\}$$

The function takes propositions and probabilistic context sets into probabilistic context sets. The essential effect of assertion is now to eliminate context credence functions. An assertion of  $p$  eliminates from the probabilistic context set all those  $Cr$  that map  $p$  to a value less than 1. The members of the probabilistic context set that results are all effectively conditionalized on  $p$ . All the  $\neg p$ -worlds are now eliminated from consideration,

<sup>26</sup>We assume here that all the c-credence functions in a probabilistic context set are defined over the same sample space, with the same  $\sigma$ -field.

as no admissible  $Cr$  function gives these worlds any nonzero c-credence. An assertion of  $p$ , to put it all another way, disallows accepting  $\neg p$  to any degree.<sup>27</sup> If we want, we can formalize the idea of taking a proposition for granted in terms of assertion: a proposition  $p$  is *taken for granted* in a probabilistic context set  $C$  just in case  $C$  is nonempty and  $\text{assert}(p, C) = C$ . If an assertion makes no change to the context, the proposition asserted is already taken for granted.

It will also help to define a more general notion of acceptance for a sentence in context. Suppose that a sentence  $\phi$  in context determines, as in Fig. 4 above, a context-change potential  $\text{ccp}$ . Then  $\phi$  is *accepted* by a probabilistic context set  $C$  just in case  $\text{ccp}(C) = C$ . It is *incompatible* with  $C$  just in case  $\text{ccp}(C) = \emptyset$ . (If  $C = \emptyset$ , call  $C$  *broken*.) Intuitively, a sentence is accepted in a context just in case uttering it won't produce any change in the context (*modulo*, as usual, the addition of the information that the utterance was just made).

What conversation will do on this model is converge, not merely on an increasingly narrow set of possible worlds, but on an increasingly narrow set of possible context credence assignments to worlds.

Now to the epistemic modals. Once again, I will take these modals to designate change functions like  $\text{assert}$ . Like assertion, I will see epistemic modal statements as imposing restrictions on the probabilistic context set, restrictions that are satisfied by eliminating context credence functions. Unlike assertion, however, the restrictions that epistemic modal statements impose will generally not eliminate possibilities from consideration altogether.

Start with 'might', 'possibly', 'probably' and 'likely'. I suggest that these modals designate the following change functions:

$$\begin{aligned} \text{might}(p, C) &= \text{possibly}(p, C) = \{Cr \in C : Cr(p) > 0\} \\ \text{probably}(p, C) &= \text{likely}(p, C) = \{Cr \in C : Cr(p) > .5\} \end{aligned}$$

The effect of saying MIGHT  $\phi$  or POSSIBLY  $\phi$  is to narrow the range of admissible context credence assignments to  $p$  by eliminating from the probabilistic context set those  $Cr$  such that  $Cr(p) = 0$ . The context that results is now incompatible with an assertion of  $\neg\phi$ . The effect of saying

<sup>27</sup>It is not difficult to imagine other ways to go. A salient alternative is to model the change induced by an assertion of  $p$  by conditionalizing with  $p$  on each  $Cr$  function in  $C$ . I am taking the current approach because it seems to me the simplest way to transform a probabilistic context set into one where every  $Cr$  assigns  $p$  degree 1 c-credence. Perhaps the more complicated conditionalizing operation will turn out to be one we need, but let us wait for the evidence favoring it. (I suspect such evidence might be found when we attempt to analyze indicative conditionals in this setting.)

PROBABLY  $\phi$  or LIKELY  $\phi$  is to narrow the range of admissible context credence assignments to  $p$  by eliminating from the probabilistic context set those  $Cr$  such that  $Cr(p) \leq .5$ . According to the context that results,  $p$  is more accepted than  $\neg p$ , as every admissible  $Cr$  assigns  $p$  the greater value. And again, the context that results is now incompatible with an assertion of  $\neg\phi$ .

This approach evades arbitrary real-value assignments to the modals, and intuitive entailment relations between sentences can now be captured. The kind of entailment relation we need, however, is a context-relative one: since the meanings of epistemically modalized sentences will be given fundamentally by context-change potentials and not by sets of truth-conditions, standard model-theoretic or possible worlds consequence won't suffice. Call the sense of 'entailment' required *context entailment*, or *c-entailment*.

A sequence of sentences  $\phi_1 \dots \phi_n$  *c-entails*  $\psi$  just in case every probabilistic context set which accepts each of the sentences  $\phi_1 \dots \phi_n$  also accepts  $\psi$ .<sup>28</sup>

Now we have it that an assertion of

It's raining,

c-entails

It's probably raining.

which itself c-entails

It might be raining.

which appears to be the intuitively correct result.

In the next three subsections, I consider how the account fares on the puzzling conjunctions that occupied us for the first half of the paper, the eavesdropping cases, and the semantics of 'must'.

## THE PUZZLING CONJUNCTIONS EXPLAINED

Start with our puzzling conjunctions. The explanation of the defectiveness of sentences like

(2) It isn't raining and it might be raining.

---

<sup>28</sup>This is basically the notion Stalnaker (1975) calls "reasonable inference". Cf. also Veltman's "validity<sub>s</sub>" in Veltman (1996).

- (3) It isn't raining and it is possibly raining.  
 (18) It isn't raining and it is probably raining.  
 (19) It isn't raining and it is likely raining.

within the current analysis should be, I hope, intuitively clear: the idea is that the conjuncts place incompatible demands on the probabilistic context set, thereby breaking it. The incompatibility between the conjuncts can be represented as follows. Let  $p$  be the proposition that it's raining, and  $\neg p$  the proposition that it isn't raining. Let 'and' designate **and**, a function from pairs of context-change potentials to context-change potentials:

$$\text{and}(\text{ccp}_1, \text{ccp}_2, C) = \text{ccp}_2(\text{ccp}_1(C))$$

Then the reason that (2), (18), and (19) are defective is that, relative to any credal context set  $C$ ,

- (2')  $\text{might}(p, (\text{assert}(\neg p, C)) = \emptyset$   
 (3')  $\text{possibly}(p, (\text{assert}(\neg p, C)) = \emptyset$   
 (18')  $\text{probably}(p, (\text{assert}(\neg p, C)) = \emptyset$   
 (19')  $\text{likely}(p, (\text{assert}(\neg p, C)) = \emptyset$

Note that the order of the conjuncts does not matter; on either ordering, the resulting context will be broken.<sup>29</sup>

If this explanation is correct, schematic renderings of these sentences like  $\ulcorner \neg\phi$  and MIGHT  $\phi \urcorner$  are strictly speaking mistaken. The correct form would be  $\ulcorner \text{ASSERT } \neg\phi$  and MIGHT  $\phi \urcorner$ , where **ASSERT** is that element (whatever it is) indicating **assert**. The schematic letter  $\phi$  here should be understood taking as its substitution class not sentences (which we have generally been saying designate, in context, context-change potentials) but rather

<sup>29</sup>This marks a difference with the otherwise similar kind of explanation Veltman's work (Veltman (1996)) makes possible for 'might'. According to his semantics for 'might' (and keeping in place our assumptions concerning **and**),  $\ulcorner \neg\phi$  and MIGHT  $\phi \urcorner$  always yields the empty set, whereas  $\ulcorner \text{MIGHT } \phi$  and  $\neg\phi \urcorner$  does not. The reason is that in the latter case the test corresponding to the modal is performed (and passed) before the  $\neg\phi$  possibilities are eliminated. To explain why this ordering of the conjuncts is defective, Veltman must appeal to some further requirement on conjunctions. (For instance, he might lay down a requirement that both conjuncts be accepted in the final context set; see the remarks on "coherence" in Groenendijk et al. (1995)).

‘sentence-radicals’<sup>30</sup>, syntactic elements which designate only propositions *sans* force.

The account of the incoherence of supposing these sentences is structurally the same, if we characterize the attitude of supposition in the same basic way that we have characterized acceptance in context. Let any given case of supposition be characterized by a set of credence functions, the ones compatible with what is being supposed. Then a sentence characterizes what is being supposed by the agent just in case the sentence is accepted by that set. Since the above sentences are only accepted by the empty set, this is a way of representing the incoherence of supposing one of them (or for that matter, instructing another person to suppose one of them).

As noted above, some find it already odd to say

(36) ? Suppose it might be raining.

A natural conjecture that these people reserve ‘suppose’ to refer to an acceptance-like attitude that does not admit of degrees, so that the modal is out of place here.

## EAVESDROPPING

What about the eavesdropping cases? When (as above) Jane rejects the MIGHT  $\phi$  claim of Sally, she is rejecting the context change that would be required of her were she to be actually sharing a discourse context with Sally. She declines, in other words, to accept the sentence.<sup>31</sup> She does so simply because she believes proposition being weakly put forward by Sally is false. Within the current framework, we can see the intuitions revealed by these cases as tracking the truth-conditional content of the overheard sentences. It’s just that the truth-conditional content of these sentences is simply the embedded proposition. The modal does not take a proposition into a new proposition; it just modifies the force with which the embedded proposition is put forward. End of story.

Almost. Upon hearing these eavesdropping cases, most people appear to think the eavesdropper does makes a warranted, true-sounding claim. But most people also appear to have the (seemingly conflicting) intuition that the eavesdroppee didn’t actually violate any kind of norm in saying what she said. This is worth remarking upon. It may be that, in addition

<sup>30</sup>The terminology comes from Wittgenstein (Wittgenstein, 1953, p.11) (“Satzradikal”) via Stenius (1967).

<sup>31</sup>Accept the sentence where?—with respect to her own context, of course. Where else?

to making distinctive changes to the context, epistemic modal claims are also subject to distinctive norms, norms different from that of assertion.

But wait—what is the norm of assertion? Better: *is* there a norm of assertion? So far I have characterized assertion as a specific kind of context change, not as an action subject to certain norms. Are there nevertheless norms attaching to the context change I have defined? Considered in isolation, it seems doubtful. This context change will be deployed in diverse circumstances, from ordinary information exchange, to remarks made under suppositions, to the interpretation the utterances of actors on a stage. If we want to attach a norm to this context change, we must restrict ourselves. The following seems plausible: let us restrict our attention to (1) communicative contexts wherein it is common ground that the discourse participants are aiming to communicate information (truths), and (2) sentences in context wherein an **assert** function takes widest scope. And let us temporarily restrict our use of “assertion” to refer to acts of uttering these kinds of sentences in those kinds of contexts. Now we can ask: is assertion subject to any norms?

Probably. After all, we are talking about cooperative agents with specific goals in conventional contexts, just sort of the place one should expect to find established norms for achieving goals. Here is not the place to defend in any detail an account of what the norm of assertion might be. I will just assume an account without argument, one which I think makes for a fruitful contrast with epistemic modal statements. I will assume the knowledge account: in order to correctly assert a proposition  $p$ , one must know  $p$ .<sup>32</sup>

Let us consider now epistemic modals claims, restricting our attention to (1) communicative contexts wherein it is common ground that the discourse participants are aiming to communicate information (truths), and (2) sentences in context wherein an epistemic modal take widest scope. Call the acts of uttering these kinds of sentences in these kinds of contexts *quasi-assertions*. (“Quasi” because although these acts involve, like assertion, restricting the possible c-credence values in a proposition, they standardly restrict the values to a lesser degree.) If knowledge is the norm of assertion, what (if anything) is the norm of quasi-assertion?

If quasi-assertion has a norm, I am not at all sure how to express it. But in order to do justice to the intuition voiced above, viz., that the eavesdroppee didn’t actually violate any kind of norm in saying what she says, we need not actually articulate the norm. We just need to suppose

---

<sup>32</sup>For a defense of this account, see Williamson (2000), chapter 11. The remarks that follow do not actually require the knowledge account to be correct for assertion; I just think the contrast that account makes possible is helpful here.

there exists the following sort of contrast:

For an assertion to be correct, the proposition asserted must be true. But for a quasi-assertion to be correct, the proposition quasi-asserted need not be true.

If the knowledge account is the right one for assertion, then truth is required for correctness. But quasi-assertion only needs, it seems, to be appropriately responsive to the evidence to be correct: the way a quasi-assertion restricts the possible *c*-credence values in a proposition *p* only needs to be appropriately responsive to the evidence about *p* available to the speaker at the context. What counts as being “appropriately responsive to the evidence” is a very dark matter. But whatever it is, all we need to suppose about it is that one can be appropriately responsive to the evidence about *p* and correctly quasi-assert *p* even when—as it turns out—*p* is false. So it appears, at least, if the intuition that eavesdroppee violated no norm is correct.

#### NEGATION AND SEMANTICS FOR ‘MUST’

Can we make sense of the traditional idea that ‘must’ is the dual of ‘might’ on this account—that **MUST**  $\phi$  is equivalent to  $\neg$ **MIGHT** $\neg\phi$ ? We can, but to do so we need to get clear on negation within the present system. It must be allowed to come in two guises: first, as the usual function from propositions into propositions, and second as a function from context-change potentials into context-change potentials. The former is needed for mundane negated sentences like ‘It’s isn’t raining’, which are best analyzed as expressing functions of the form **assert**( $\neg p$ ). The latter is needed for negation when it takes scope over items expressing context-change potentials, like epistemic modal clauses.<sup>33</sup> This function we define in the style of Heim (1983):

$$\text{not}(\text{ccp}(C)) = C \setminus \text{ccp}(C)$$

Now we have it that

$$\text{not}(\text{might}(\neg p, C)) = C \setminus \text{might}(\neg p, C)$$

and since

$$C \setminus \text{might}(\neg p, C) = C \setminus \{Cr \in C : Cr(\neg p) > 0\}$$

---

<sup>33</sup>I am not exactly sure that negation can naturally take scope over epistemic modals; most of the examples I know involve the rather artificial “It is not the case that” locution. But supposing it can, what follows is what, I think, we’d need to say about it.

and  $C \setminus \{Cr \in C : Cr(\neg p) > 0\}$  is just  $\{Cr \in C : Cr(p) = 1\}$ , the result is that, if ‘must’ is the dual of ‘might’, the change function it expresses may be given as:

$$\text{must}(p, C) = \{Cr \in C : Cr(p) = 1\}$$

This should look familiar. It is the same change function as **assert**.

Does this make the right predictions? It is widely thought, after all, that ‘must’ conveys something “weaker” than the corresponding assertion without ‘must’. Wouldn’t it be better, then, to have a change function that invokes a change weaker than **assert**? Perhaps something more like

$$(?) \text{ must}(p, C) = \{Cr \in C : Cr(p) > .9\}$$

or, more vaguely,

$$(?) \text{ must}(p, C) = \{Cr \in C : Cr(p) \text{ is very high}\}$$

This would make ‘must’ akin to something like ‘very likely’ or ‘almost certainly’, on the most natural way of modeling these latter in the current system.

But this makes the wrong predictions. Suppose you buy a lottery ticket. I may say to you

(37) You might be the winner, but you very likely are not the winner.

(38) You might be the winner, but you are almost certainty not the winner.

but I cannot say

(39) # You might be the winner, but you must not be the winner.

So ‘must’ is not equivalent to something like ‘very likely’ or ‘almost certainly’. If **MUST**  $\phi$  is accepted in a context, the proposition expressed by  $\phi$  is taken for granted in that context.

One idea is locate the felt difference between **ASSERT**  $\phi$  and **MUST**  $\phi$  outside the formal semantics, and in the norms by which the acts of uttering these sentences are evaluated. The sentences invoke the same context changes, but the acts of uttering these sentences are subject to different conditions of correctness. **ASSERT**  $\phi$  requires knowledge of the proposition expressed by  $\phi$ . **MUST**  $\phi$  requires less: some kind of appropriate responsiveness to the evidence falling short of knowledge.

## 12

This has only been a brief sketch of a draft of an I.O.U. of an outline of one idea about how to spell out a force analysis of epistemic modals. Obviously, the work is very incomplete—we don't yet have all the details on the table. It remains to show how epistemic modals interact with conditionals, attitude verbs, and tenses on this approach. It remains to be settled what exactly represents **assert**, and whether this change function might not be better treated as a conditionalizing operation. More needs to be said about how best to understand the attitude of acceptance. Adjustments will undoubtedly need to be made. I hope this beginning has looked promising enough to motivate interest in further investigation.<sup>34</sup>

## References

- DeRose, K. "Epistemic Possibilities". *Philosophical Review*, 100(4), (1991), 581–605.
- Edgington, D. "Do conditionals have truth-conditions?" *Critica*, XVIII(52), (1986), 3–30.
- Egan, A. "Epistemic Modals, Relativism, and Assertion". In J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin, eds., *New Work on Modality* (MIT Working Papers in Linguistics, 2005).
- Egan, A., J. Hawthorne, and B. Weatherson. "Epistemic modals in context". In G. Preyer and G. Peter, eds., *Contextualism in Philosophy* (Oxford University Press, 2005).
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge* (Cambridge, MA: MIT Press, 1995).
- von Fintel, K. and A. Gillies. "'Might' Made Right" (2004). Unpublished handout.
- von Fintel, K. and S. Iatridou. "Epistemic Containment". *Linguistic Inquiry*, 34(2), (2003), 173–98.
- Groenendijk, J., M. Stokhof, and F. Veltman. "Coreference and Modality". In S. Lappin, ed., *Handbook of Contemporary Semantic Theory* (Blackwell, 1995).
- Hacking, I. "Possibility". *Philosophical Review*, 76(2), (1967), 143–68.

---

<sup>34</sup>Manifestly, this is work in progress. For helpful discussion, thanks to Stephen Yablo, Robert Stalnaker, Andy Egan, Selim Berker, Danny Fox, Ephraim Glick, Ned Hall, and especially Dilip Ninan.

- Heim, I. "On the projection problem for presuppositions". In M. Barlow, D. Flickinger, and M. Wescoat, eds., *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics*, 114–25 (Stanford University, 1983).
- Jeffrey, R. "Indefinite Probability Judgment: A Reply to Levi". *Philosophy of Science*, 54(4), (1987), 586–591.
- Jevons, W. S. *Elementary Lessons in Logic: Deductive and Inductive* (London and New York: Macmillan and Co., 1878).
- Kaplan, D. "Demonstratives". In J. Almog, J. Perry, and H. Wettstein, eds., *Themes from Kaplan*, 481–563 (Oxford University Press, 1989).
- Kaplan, M. *Decision Theory as Philosophy* (Cambridge University Press, 1996).
- Kneale, W. *Probability and Induction* (Oxford University Press, 1949).
- Kratzer, A. "The notional category of modality". In H.-J. Eikmeyer and H. Rieser, eds., *Words, Worlds, and Contexts, New Approaches to Word Semantics*, 38–74 (Berlin: Walter de Gruyter, 1981).
- Lewis, D. "Attitude De dicto and de se". *Philosophical Review*, 88(4), (1979), 513–43.
- . "Index, Context, and Content". In S. Kanger and S. Ohman, eds., *Philosophy and Grammar*, 79–100 (Holland: D. Reidel, 1980).
- MacFarlane, J. "Epistemic modalities and relative truth" (2003). Unpublished.
- Moore, G. E. *Commonplace Book 1919-1953* (London: George, Allen, and Unwin, 1962).
- Stalnaker, R. "Pragmatics". *Synthese*, 22, (1970), 272–89.
- . "Indicative Conditionals". *Philosophia*, 5.
- . "Assertion". In P. Cole, ed., *Syntax and Semantics 9: Pragmatics* (Academic Press, 1978).
- . "A Defense of Conditional Excluded Middle". In W. Harper, R. Stalnaker, and G. Pearce, eds., *Ifs*, 87–104 (Dordrecht: Reidel, 1981).
- . *Inquiry* (Cambridge: MIT Press, 1984).
- Stenius, E. "Mood and Language-Game". *Synthese*, 17, (1967), 254–74.
- Thomson, W. *An Outline of the Necessary Laws of Thought: A Treatise on Pure and Applied Logic* (New York: Sheldon and Company, 1877). Fourth edition.

Toulmin, S. "Probability". In A. Flew, ed., *Essays in Conceptual Analysis*, 157–91 (New York, NY: MacMillan, 1956).

Veltman, F. "Defaults in Update Semantics". *Journal of Philosophical Logic*, 25(3), (1996), 221–61.

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

Wittgenstein, L. *Philosophical Investigations* (Blackwell, 1953), 2nd edition. Translated by G. E. M. Anscombe.