

## Eavesdroppers and Epistemic Modals<sup>1</sup>

A simple and initially quite natural account of a sentence governed by an epistemic modal says that its truth or falsity at a context is determined by the knowledge possessed by the speaker at that context: ‘Might P’ (and ‘It could be that P’) is true at a context iff P is compatible with what the speaker knows. ‘Must P’ (and ‘It has to be that P’) is true at a context iff P is entailed by what the speaker knows.<sup>2,3</sup> Some kind of fine tuning is no doubt needed here: Perhaps the relevant body of knowledge is sometimes not just the speaker’s but the combined knowledge of all parties to the conversation in which the speaker is engaged. Perhaps ‘compatible’ needs to be replaced by ‘not obviously incompatible’.<sup>4</sup> And perhaps we need to supplement the account with a pragmatics that explains why it is unacceptable to say either ‘Might P’ or ‘Must P’ when P is part of what one has to go on. Moreover, for some uses of past tense ‘might’ claims, the relevant body of knowledge is that which the speaker used to possess (as when one says that it might have rained in a context where one is explaining why one took an umbrella).<sup>5</sup> But these points are compatible with the view that the simple account is very much on the right track.

However, there are some data that suggest that there is something seriously wrong about the simple account. The following case is representative.

EAVESDROPPER: Someone is on the way to the grocery store. I hear her say: ‘Susan might be at the store. I could run into her.’ No party to the conversation that I am listening in on knows that Susan is on vacation. But I know that she is. Despite the fact that it is compatible with what the conversants know that Susan is in the store and that the speaker will run into her, I am inclined to judge the speaker’s modal judgments to be incorrect.

---

<sup>1</sup> Thanks to Kai von Fintel, Timothy Williamson, Dorothy Edgington and audiences at St Andrews and the 2007 Sofia conference in Mexico for helpful feedback.

<sup>2</sup> I assume in this paper that sentences involving the epistemic modals ‘might’, ‘must’, ‘could’ and ‘has to’ have truth conditions (at contexts). There is a good prima facie evidence for this assumption. For example, these sentences -- unlike ones of the form ‘Perhaps P’ -- embed smoothly in the antecedents of conditionals, a position that will not tolerate the more obvious examples of non-truth-apt sentences.

<sup>3</sup> These formulations suggest that ‘might’ and ‘must’ are sentential operators. I by no means wish to commit myself to this. Nor, in using the expression ‘epistemic modal’ do I wish to imply that there is a lexical ambiguity between purely epistemic and other uses of ‘might’. (For a story that militates against any claim of ambiguity, see Angelika Kratzer’s ‘What ‘must’ and ‘can’ must and can mean in ‘Linguistics and Philosophy,’ January 1977, 337-355 and ‘The Notional Category of Modality,’ in H.J. Eikmeyer and H. Reiser, eds., *Words, Worlds and Contexts*, Springer, 1981, 38-74 ).

<sup>4</sup> For relevant discussion, see Keith DeRose, ‘Epistemic Possibilities,’ *The Philosophical Review*, 1991, 581-695.

<sup>5</sup> Interestingly, ‘It might have rained’ has this shifty reading, but ‘It must have rained’ does not. I defer to linguistics to explain this. For relevant discussion, see Cleo Condoravdi ‘Temporal Interpretation of Modals: Modals for the Present and for the Past,’ in *The Construction of Meaning*, edited by David I. Beaver, Luis D. Casillas Martínez, Bardy Z. Clark, & Stefan Kaufmann. Stanford: CSLI, 2002, and Tim Stowell, ‘Tense and Modals,’ in *The Syntax of Time*, edited by Jacqueline Guéron & Jacqueline Lecarme. MIT, 2004, 621–636.

A good deal of recent discussion of epistemic modals has been directed towards understanding how to accommodate data such as this.<sup>6</sup> In what follows I shall examine and dismiss one strategy for accommodation (that I was once unwisely sympathetic to)<sup>7</sup>, and then describe a much more promising strategy.

### *A Bad Strategy: Relativism*

Contemporary relativist discussions tend to proceed along the following lines.<sup>8</sup>

First, the relativist takes heart from the tendency in mainstream philosophy of language to claim that propositions – the semantic values of that-clauses -- are true only relative to certain parameters of evaluation. In particular, orthodoxy says that truth is relative to a world of evaluation, so that the same proposition can be true relative to one world and false relative to another.<sup>9</sup> (Moreover, an important minority that includes David Kaplan says that truth is relative to a time of evaluation, so that the same proposition can be true at one time and false at another).<sup>10</sup>

Second, the relativist suggests that we allow in extra parameters of evaluation that, roughly, encode the standard of evaluation of a judge. Applied to ethics, the idea will be that different communities utter the same proposition by ‘Spitting is bad’ even if their standards are different, but that the truth of the proposition expressed is relative not only to worlds and times but also to standards. Applied to epistemic modals, the relativist maintains that (i) different communities utter the same proposition by ‘It might rain’ even if they have access to different bodies of evidence, but that (ii) the truth of the proposition expressed is relative not only to worlds and times but also to bodies of evidence.

Third, the relativist offers a distinctive account of how the evaluation of someone else’s utterances proceeds. The main idea here is that when someone evaluates some utterance, the appropriate answer to ‘Is that utterance true?’ is not given by the standards of the utterer but by the standards of the evaluator. Let us allow ourselves a ‘true at’ locution in the metalanguage, and say that a proposition is true at a world, time, standards triple iff the proposition is true relative to the world time and standards of that triple. Let

---

<sup>6</sup> See Egan, Hawthorne and Weatherson, ‘Epistemic Modals in Context’, in Gerhard Preyer and Georg Peter (eds) *Contextualism in Philosophy*, OUP 2005, pp. 131-69, Andy Egan ‘Relativism, Epistemic Modals, and Assertion,’ (forthcoming in *Philosophical Studies*), John MacFarlane ‘Epistemic Modals are Assessment-Sensitive,’ (draft), and Kai von Fintel and Anthony Gillies ‘CIA Leaks,’ (draft).

<sup>7</sup> Thus I am no longer happy with the proposal that gets most favorable treatment in ‘Epistemic Modals in Context,’ (Egan, Hawthorne and Weatherson), in Gerhard Preyer and Georg Peter (eds) *Contextualism in Philosophy*, OUP 2005, 131-69.

<sup>8</sup> For more details, see John MacFarlane ‘Making Sense of Relative Truth,’ *Proceedings of the Aristotelian Society* 105 (2005), 321-39, ‘Relativism and Disagreement’, *Philosophical Studies* 132 (2007), 17-31, Egan, Hawthorne and Weatherson, ‘Epistemic Modals in Context,’ (op. cit), and Peter Lasnik ‘Context Dependence, Disagreement and Predicates of Personal Taste,’ *Linguistics and Philosophy* (2005), December 2005, 643-686.

<sup>9</sup> For what its worth, my own best guess is that truth and falsity (as opposed to certain more gerrymandered relations that one might choose to talk about for certain purposes of convenience) are, pace orthodoxy, straightforwardly monadic.

<sup>10</sup> See David Kaplan, ‘Demonstratives,’ J. Almog, J. Perry and H. Wettstein (eds.): *Themes from Kaplan*. Oxford University Press, 1989:481-563

us also introduce a ‘true for’ locution, governed by the rule that an utterance made by X is true for X just in case it is true relative to the world, time and standards *operative* for X, which in the simple case is the world and time of the utterance and the standards of the utterer. We can now express the second idea by allowing a monadic truth predicate of utterances into the object language and crafting matters so that the utterance ‘Utterance U by X is true’ is true for Y iff the proposition expressed by X using U is true at a triple that consists of the world and time operative for X when X uttered U *and the standards operative for Y*. (Similarly, *mutatis mutandis*, for ‘is false’.)

I have standards S1 and you have standards S2 and you say ‘Recycling is obligatory’. Suppose the proposition expressed is false at the triple consisting of your world time and S2, but true at the triple consisting of your world, time and S1. The claim ‘That utterance is true’ (said of your speech) is true for me and thus felicitously assertable by me, even though the proposition you expressed was false relative to the standards that you were operating with.

How is this all supposed to help with Eavesdropper? Let the relevant parameter for epistemic modals be a body of knowledge parameter. And suppose that this parameter interacts with utterance truth in the way just described. Suppose someone utters ‘Susan might be in the store’ when I, the eavesdropper, know that Susan is not in the store. Then, by the relativist rules of the game, I can say ‘That statement is false’ and say something that is true for me. This then is offered as the proper explanation of the data.

To have even *prima facie* plausibility, this story has to be complicated further so as to allow that the body of knowledge that is operative for me when I make an assertion does not always include my own.<sup>11</sup> If Suzy hates me and gets on the bus and hides, then, with the situation filled out appropriately, I can say ‘Suzy is hiding because I might be on that bus’. Here the body of knowledge that is operative is not my knowledge but Suzy’s. The body of knowledge I intend to be relevant and that does not always include my own, and this needs to be reflecting in the relevant notion of ‘operative standards’.

Not unrelatedly, and as Kai von Fintel and Anthony Gillies have emphasised, the eavesdropping data is considerably exaggerated. Given that it is easy to get into the frame of mind where I offer ‘I might be on the bus’ as an acceptable explanation of why Susy is hiding, it is also easy to get into the frame of mind where I say that ‘That’s true’ when Susy hides and says ‘John might be on this bus’.<sup>12</sup> This generalizes to other many other examples. Indeed, once the relativist concedes that the operative body of evidence for an assertion need not include the evidence possessed by the assertor, she should *not* predict that we invariably respond with an intuitive verdict of ‘False!’ in an eavesdropper context in which we have evidence that excludes the proposition that the subject claims might be true. Acknowledging the point above about operative standards, the relativist must retreat to the view that we sometimes and only sometimes respond with an intuitive verdict of falsehood in eavesdropping cases. (Of course she could still try to run an anti-

---

<sup>11</sup> Cf. Egan, Hawthorne and Weatherson, p. 136. We tried to explain away much of the data of this sort by saying that the relevant claims are not simple assertions but rather verbalizations of someone else’s thoughts (or the thought they should be having). This now seems a bit strained to me (rather in the way that it would be strained to say that the claim ‘Bill turned left’ is not a simple assertion in the case that the operative orientation is not my own but Bill’s).

<sup>12</sup> See ‘CIA Leaks,’ (draft).

contextualist argument by claiming that even this data cannot be accommodated by the standard contextualist.)

My complaint against the baroque semantics that emerges from relativist discussions does not concern its internal coherence. It is rather that its proponents (including my former self) typically pay insufficient attention to the shape of the data, and are insufficiently imaginative and patient about trying to explain the data within a more sober framework.<sup>13</sup> In the next section I shall attempt to provide a non-relativist account of this sort. For now, let me underscore a couple of particular troubles that afflict relativist proposals. (What follows is largely a restatement of points raised in von Fintel and Gillies' 'CIA Leaks'.)

Note first, and most obviously, that the relativist strategy does not yield plausible results when applied to 'must'. Suppose I don't know whether Susan is in the store and I don't know whether Frank knows this. I eavesdrop on Frank who says 'Susan must be in the store'. I think there is very little intuitive pull to the reaction 'That's false'. Instead I react with the thought that I don't know whether Frank said something true since I don't know whether he knows where Susan is (unless, that is, I am willing to trust his testimony and so take his speech as showing that he does know where she is). The relativist does not predict this. After all, the relativist account would lead one to expect contexts in which the eavesdropper's epistemic environment is operative for 'Frank's utterance was false' and thus where that assertion is felicitous. And whenever the epistemic context of the speaker was operative, then by the above rules of the game, 'Frank's utterance was false' would be true for the eavesdropper.<sup>14 15</sup>

Second, there is a contrast in the data that the relativist has overlooked. Suppose I know where Frank's keys are. He comes and says 'Lets see. They might be in the bathroom and they might be in the kitchen'. I do not find myself with any strong inclination at all to reckon Frank's utterance as false, far less inclination that in the eavesdropper case described earlier.<sup>16</sup> In this setting it seems pretty clear that insofar as Frank's knowledge leaves both key locations open, his speech is true. One thing that needs explaining, then, is why we are intuitively inclined to use Frank's body of evidence as the arbiter of truth in this case, but are much more ambivalent about using the utterer's body of evidence as the arbiter of truth in the grocery store case. As a further clue, consider a version of the Suzy-on-the-bus case where I do not in fact ever travel on buses, even though Suzy does not know this. Now in this case too it is true that for all Suzy knows I am on the bus.. But in this case I am far less including to offer 'I might be on that bus' as an explanation of Suzy's behaviour, and relatedly, am far less inclined to judge Suzy assertion as correct. The relativist framework I have outlined has no way of explaining why my never travelling on buses might make a difference to the intuitive verdict: For whether on not I ever travel on buses, the Suzy case will be one in which I

---

<sup>13</sup> I think the same lesson applies to the currently trendy relativisms about predicates of personal taste, but I shall not pursue that here.

<sup>14</sup> Also, for many cases, I have a hard time getting into the frame of mind where I judge 'That is true' when a speaker says 'It must be that P' in a situation where I know that P and the speaker mistakenly thinks he knows. (Suppose someone thinks they have checked every room but the kitchen for the keys, but in fact have not checked the living room. The person says 'The keys must be in the kitchen' and they in fact are....).

<sup>15</sup> For a related worry about 'must', see von Fintel and Gillies, 'CIA Leaks', p. 12.

<sup>16</sup> Von Fintel and Gillies make this point in 'CIA Leaks', section seven.

know I am not on the bus and for all Suzy knows I am. Whether the operative body of evidence is mine or Suzy's, the fact that I in fact never travel on buses would not count as relevant to my evaluation of Suzy's utterance.

### *Epistemic Modals and our Concept of Danger*<sup>17</sup>

In what follows, I shall sketch an approach that is more promising than relativism, and which, in particular, does rather better with the complexities of the eavesdropping data. Let me acknowledge at the outset, however, that there may be causes of the eavesdropping data other than those that I am about to focus on. Notably, there are settings in which one runs roughshod over differences in meaning that are due to differing standards on account of the fact that one thinks one's own standards are in some sense superior. Thus suppose I am a big spaghetti eater and Roger is not. Roger gets served by what is by my standards a moderate amount and says 'That's a lot of spaghetti!' I may react by saying 'That's not a lot of spaghetti. He needs to try to servings and Julio's and then he'll know what a lot of spaghetti is.' In this setting I talk as if Roger has made a mistake. But I may not be very serious in this conviction. What I am mainly trying to do is to make sure that my standards for 'a lot' govern my current conversation and moreover wish to communicate my superiority on at least one dimension by doing so. It would be ridiculous to use data like this as a basis for constructing a baroque relativist semantics for 'a lot'. What we need instead is some socio-pragmatic account of why we often impute semantic uniformity in only a half-serious way as a device to achieving this or that communicative goal. (Consider similarly someone from Arizona visiting New England who criticizes a Bostonian who says 'Its hot' on a 97 degree day. Again, this is only a half-serious criticism – fuelled by a half-serious pretense that Arizonian are more or less unique in being correctly calibrated for heat detection. It would again be silly to think that data such as this is a compelling basis for relativism about 'hot'.) At least some of the eavesdropping data for epistemic modals may have its origin in similar mechanisms.<sup>18</sup> That said, I shall turn to what I take to be a more fundamental source of the eavesdropping data when it comes to modals.

As a prelude to my main proposal, some groundclearing is in order regarding the concepts of chance at play in our ordinary conceptual scheme. Many philosophers (especially those brought up on David Lewis) find it natural to think there are basically two notions of chance<sup>19</sup>: On the one hand there are epistemic notions of chance, that are tied to the epistemic state of the agent; on the other hand there is the notion of objective

---

<sup>17</sup> I was helped here by work in progress by Antony Eagle.

<sup>18</sup> A further possible source of the eavesdropping data is discussed by von Fintel and Gillies ('CIA Leaks', draft): when an epistemic modal is asserted for public consumption, the size of the group whose body of knowledge is relevant extends far beyond the conversational participants. They briefly go on to suggest an interesting twist on this idea: in many cases it is a vague matter how far the relevant group extends ( the speaker can then precisify when challenged).

<sup>19</sup> I note in passing that those familiar with statistical mechanics will not be so disposed to a taxonomy of this sort.

chance, which is future directed (the objective chance that P for any present or past proposition is 1 or zero), and is determined to be intermediate between zero and 1 not by the epistemic state of the agent but rather by whether the laws of nature plus the state of the world leave a certain future open.

This is far too impoverished a taxonomy. Take, for example, our ordinary notion of danger<sup>20</sup> (and its optimistic counterpart, hope). It fits neither the epistemic chance model, nor the objective chance model outlined above. Suppose I see what appears to be a person tottering at the edge of a cliff. Suppose in fact my eyes are deceiving me. The person is nowhere near the cliff. If I think to myself ‘The person is in danger of falling off,’ I am saying something false. But it is clear that for all I know the person will fall off – so an epistemic gloss will not account for the falsity of my utterance. Note that I do not seem to be committing myself to a failure of determinism in making the claim that the person is in danger. Consider another example. Suppose I find out that someone has a very dangerous disease. It would be perfectly felicitous to say ‘The person is in danger of dying. Whether they end up surviving or not will depend on all sorts of subtle factors including the strength of their immune system and so on.’ Here I am acknowledging that the finer details may secure their survival even while asserting there is a danger they will die.

This notion of danger does not merely apply to the future. Suppose I have just brought home a washing machine. I am told that a recall has been put out for that brand owing to the fact that many are defective. I say ‘There is a danger that this is defective’. This is not merely epistemic modality. If the recall is in for a different brand and this is in fact a very reliably made brand, then the judgment is intuitively incorrect even it is true that for all I know, the brand is defective. Similarly, if I believe that someone is in danger of having a certain disease because I think they have been to a country where that disease was rife, then my judgment is intuitively true if they have been to that country but in fact did not acquire the disease, but intuitively false if they did not in fact visit a country where the disease was going around.

The following schematic semantics will serve our purposes here. Let a case be a centred world (in Lewis’ sense). Let us suppose some cases are close to a given case. X is at t1 in danger of Fing at t2 iff some centred worlds that are close to the actual case centred on X at t1 are worlds where x Fs at t2.<sup>21 22</sup> As with counterfactuals, closeness is a

---

<sup>20</sup> What I have to say applies mutatis mutandis to its optimistic counterpart, hope. (I have in mind the use of ‘hope’ that is not straightforwardly psychological: ‘I wonder whether he has a hope of surviving’)

<sup>21</sup> We could perhaps adjust ‘some’ to ‘sufficiently large proportion’ – where ‘sufficiently large’ is context sensitive. We will then need to postulate some measure over centred worlds to handle cases where there are uncountably many close ones

<sup>22</sup> One could also try to build a toy semantics using the standard apparatus of statistical mechanics. The approach that I have in mind assumes the measure on phase space that is crucial to calculating statistical mechanical chances (where each point in phase space is a possible combination of positions and velocities for the particles in a world, and the measure is the so-called Liouville measure.) A given body of information – a ‘macrostate’ will determine a region of the phase space. ‘There is a danger that P’ is true at a context C, if there is no sub-region of the relevant macrostate (or, perhaps, no non-negligible sub-region) where P, where the relevant macrostate is delimited by the relevant facts – in the sense discussed in the text -- at C. Of course this is only as precise as ‘relevant fact’ is precise. (The choice between ‘no’ and ‘no non-negligible’ will depend, inter-alia, upon whether one allows the proposition that there will be no quasi-miraculous developments as a relevant fact. Without some such constraint one faces the problem that pretty

somewhat context sensitive affair. But what we do know from the above considerations is that there will often be epistemically accessible cases (cases which for all the speaker knows are actual), which nevertheless do not count as close in the relevant sense. In a context there will be a rough and ready standard of relevant similarity, albeit vague. In the washing machine case, a world where the washing machine is made by a different brand is (even if possible) not close. But there are close worlds where some of the fine details of its construction – ones that vary across tokens of that brand – are different to those that are actual. In the cliff case, worlds in which there are subtle variations in the wind conditions count as close. But worlds where it is quite a long way from where it actually is do not. The context dependence and sloppiness of these judgments – in effect lack of discipline about implicit standards of relevant similarity – make systematic evaluation of danger judgments frustrating, and predicts instability in evaluations of truth value. But such judgments are nevertheless an important part of our conceptual scheme – one which populates the world with hopes and dangers as we try to make practical headway through life.

I propose that ‘might’ claims sometimes have a purely epistemic use – ones that are merely tied to the knowledge set of the agent (or perhaps agent plus interlocutors), and a more objective – ‘danger’ – use, that has no straightforward analysis in terms of bodies of knowledge. It may be that no individual at a world knows that a certain chair is in fact a long way from the edge of the cliff and hence the sum total of what everyone in the world knows leaves it open that the chair will fall off the cliff. But still, the claim ‘There is no danger of the chair falling over’ is true in the mouth of individuals at that world. When someone in the earlier example says ‘Lets see. The keys might be in the kitchen and they might be in the living room,’ it is natural to hear that as a purely epistemic use. By contrast, if a surgeon says ‘Lets try to find out whether this patient might be about to die’, the use is danger-theoretic.

Part of why our intuitions about particular cases have slippage is because it is not clear whether the epistemic or danger-theoretic use is in play. Suppose a coin is in fact double-headed and is tossed. Someone says ‘It might be showing tails’. Informants in the know tend, in my experience, to fluctuate in their assignments of truth value. At least part of the explanation of this is that some take it as merely epistemic while others do not. Certain follow up speeches by the speaker solidify intuitions dramatically. Suppose, for example, the speaker says ‘Granted it may be a two headed coin. But it may also not be. So it may be showing tails.’ Then it is overwhelmingly natural to hear the initial claim as true. Relatedly, if the original speech was made in a context in which the person was deciding whether to bet heavily on its showing tails at not especially good odds, the judgment that the speech was correct would seem particularly natural. If, meanwhile, a ‘danger’ use is in play it is natural to take the speech as incorrect. After all if the speaker had said ‘The coin is in danger of coming up tails’, he would be natural to hear him as having made a mistake.<sup>23</sup> Interestingly, ‘must’ is less liable to be used non-epistemically. To my ear, if someone said ‘The coin must be showing heads’ in the double-headed coin case, that judgment sounds incorrect. (I don’t think our judgment here is merely that it is

---

much any macrostate will have a negligible subregion corresponding to pretty much any fantastically bizarre outcome that one could imagine.)

<sup>23</sup> Of course the ‘danger’ talk is only natural in a case where the relevant outcome is bad, so imagine there was a bad consequence for tails.

unassertable.) Of course ‘It is certain that’ does have the relevant non-epistemic reading.<sup>24</sup>

Clearly, our practice of populating the world with hopes and dangers is tied to practical reasoning. We recognize that what will count as a danger from one perspective won’t count as a danger from another. For someone standing at the top of the stairs in the dodgy neighbourhood, there is a danger that the cellar will contain a criminal. But if I know there is no criminal in that cellar, there is no danger as far as I am concerned. This agent-sensitivity does not however mean that the presence or absence of a danger can be analysed in terms of the epistemic state of the agent. When Suzy hides I recognize that there is a genuine danger for her that I am on the bus. But in a case where I in fact never travel on buses I will think that even for Suzy the danger is illusory.

Back to the eavesdropping data. Someone says one Sunday morning ‘Susan might be in the local grocery store’. The account predicts that if the purely epistemic use is at work we won’t react with a ‘false!’ verdict. Suppose the speech was continued in a way that made it look purely epistemic. Suppose the person continued ‘Granted, she may have moved out of town. But she may still be around here. Really I have no idea. So she *may* be in grocery store.’ I think the eavesdropper would not be inclined to contest any of the elements of this speech. But the account I have proposed also predicts a use which is more or less tantamount to ‘There is a danger that Susan is in the grocery store’. Here the truth conditions are not given by what the speaker knows, nor even by what some larger community knows. It turns on whether there is a relevantly similar situation to the actual world where Susan is in the grocery store. If in fact she has moved from town long ago, the danger is illusory. But if she often goes to the grocery store on a Sunday morning, then it is not. Granted the line between real and illusory dangers is extremely vague – but that is not a problem with the account provided here.<sup>25</sup> It is enough for my purposes if there is a danger-theoretic use of modals that cannot be analyzed in terms of the epistemic state of groups of agents, and that this generates readings that explain much of the eavesdropping data with which we began.

---

<sup>24</sup> Reiterating an earlier theme. I think it is wrong to think that ‘It is certain that P’ is true iff the current state of the world plus the laws of nature entail P. One can say ‘It is not certain that such and such will happen’ and be right just so long as there is a relevantly similar situation which eventuates in such and such not happening.

<sup>25</sup> Granted, some will think that danger talk is so lacking in discipline that it doesn’t really have truth conditions. Reactions of that sort connect with larger methodological issues that I can’t address here.