3

'Might' Made Right

Kai von Fintel and Anthony S. Gillies

1. Introduction

The simplest story about modals—*might*, *must*, *possibly*, *necessary*, *have to*, *can*, *ought to*, *presumably*, *likelier*, and the rest—is also the canon: modals are context-dependent quantifiers over a domain of possibilities. Different flavors of modality correspond to quantification over different domains of possibilities. Logical modalities quantify over all the possibilities there are, physical modalities over possibilities compatible with the laws of physics, deontic modalities over possibilities compatible with what ought to be. And epistemic modals—and in particular epistemic *might* and *must*, the stars of the show here—quantify over possibilities compatible with *what is known*.

But fixing a flavor of modality need not always fix what the modal says since context can still figure prominently in determining the domain. That seems right for epistemic modals: in one context it is what Holmes knows that is relevant, in another it is the information Watson has at hand that counts. The modals quantify over possibilities compatible with the information at hand—they quantify over what's not ruled out by a relevant *information state*. Just whose information state is something decided by the context.

That is the canon. It is simple and elegant but not quite right. It says that bare[1] epistemic modals quantify over the information available to a contextually relevant group. The context decides the group (and perhaps the standards by which they know) and thereby the domain the modals quantify over. But—as we will argue—in

This paper has been with us—in the form of notes, handouts, and talks—since 2003. That is a long time, and we have gotten attached to it. It has even done us the favor of reproducing a time or two, spinning off bits and pieces that grew up faster than it did (von Fintel and Gillies 2007, 2008). While we're happy to have finally got it off our to-do list, we will miss it. For comments, we would like to thank Josh Dever, Angelika Kratzer, Chris Potts, and audiences at the University of Osnabrück, at the Workshop on (In)determinacy of Meaning in Cologne at the annual meeting of the German Linguistics Society, at the University of Texas at Austin (twice), and in a mini-seminar on epistemic modals and conditionals at the University of Frankfurt. We are also indebted to an anonymous reviewer for Oxford University Press.

[1] "Bare" epistemic modals because they occur without accompanying restrictors.

perfectly ordinary conversations this story gets off track, making bare epistemic modals either too hard to reasonably assert or by making their uptake or dissent mysterious.

Of course, the canon is not the only game in town. And the kind of context-dependence embodied by it has recently gone in for some rough treatment. CIA agents—those arguing that epistemic modals only get assigned truth values relative to **c**ontexts of utterance, **i**ndices of evaluation, and (the new wrinkle) points of **a**ssessment—have been trying to incite a revolution against it.[2] (Relativism is the new black.) And that kind of story is meant to be better suited for dealing with some of the puzzling behavior of epistemic modals, including we'll assume the problem we raise here. It is not a state secret that we have our doubts about whether CIA theories hold water (see e.g. von Fintel and Gillies 2008). Still, if nothing else could solve the problem we raise for the canon, we would have to reconsider.

But, we will argue, there are other options and so we need not yet reconsider. The problem with the canon isn't that it has an impoverished set of superscripts. The problem is that it is weighed down by the fiction that there is a determinate context in which an utterance of a bare epistemic modal is issued. We will sketch a story about the pragmatics and conversational dynamics of these modals that goes a bit like this. Bare epistemic modals quantify over the information available to a contextually relevant group of investigators. That part the canon gets right. But, given a context, there are multiple ways of drawing the group boundaries. And "the context" often does not decide which of these is to be preferred. Thus, it is indeterminate just which group—just which aggregated information state—is quantified over by such bare epistemic modals. That part the canon gets wrong.

We need to be more precise about the problem we see for the canon. But first we want to sketch our preferred version of it (Section 2). Then, after raising the problem as we see it (Section 3), we will turn to exploring how a story that exploits this contextual indeterminacy can explain what needs explaining without departing from what's right about the canon and so without positing anything semantically peculiar about bare occurrences of epistemic modals (Sections 4-6).

2. The Canon

We will assume a generic logical form for modals along these lines:

(1) MODAL$(B)(\varphi)$

The interpretation of MODAL is just some quantifier Q—maybe first-order definable, maybe a fancier generalized quantifier—supplying the relevant force of the modal.[3]

[2] Some recent CIA agents: Egan (2007); Egan et al. (2007); MacFarlane (this volume); Stephenson (2007a,b).

[3] More generally, the interpretation of MODAL is a quantifier Q_D over domain D, the value of D being a function of context or the preceding discourse and the modal saying that Q_D of the possibilities in B are possibilities in which the prejacent is true. Putting things this way is useful for dealing with modal subordination, but since such issues aren't our focus here, we suppress D.

The first argument *B*—the *modal base*—determines the restriction on the domain of worlds over which the modal in question quantifies. The second argument is the *prejacent*—the sentence that the modal claim says holds in *Q* of the possibilities in *B*. Vary the quantificational force, or vary the domain of possibilities determined by *B*, and we get different modals.[4]

One way (the value of) *B* gets determined is as the denotation of restricting phrases like *in view of*.[5] Some examples:

(2) a. In view of what the laws are, cars cannot park on the Turnpike.
b. In view of the information available to Holmes, the gardener might have done it.
c. In view of the preferences Sally has, she has to take the early train.

The modals quantify over the possibilities compatible with the laws, the information Holmes has, and the satisfaction of Sally's preferences. And in each case the modal base is (plausibly) just the denotation of the relevant restricting *in view of* phrase that occurs earlier in the sentence.

But sometimes there is no restricting phrase; then *B* must be determined in some other way. Some examples:

(3) a. Sally presumably will not make it to the meeting.
b. Alex must be hungry.
c. Your keys might be on the desk.

These are *bare modals* since there is no restrictor explicit. The canon says that what helps determine the value of *B* for bare modals is "the context." Our focus here is on bare *epistemic* modals (BEMS): bare modals where the modal base *B* supplied by "the context" is epistemic and thus (the denotation of) *B* is the set of worlds compatible with the relevant information state. To repeat the rough gloss of the canon: such modals are quantifiers over possibilities compatible with the available evidence or the information at hand or the relevant information state—that is what makes them *epistemic* modals—saying that all/some/most/just the right such possibilities are possibilities in which the prejacent is true.

So *might* and *must* are quantifiers over information states, true at a context–index pair just in case the prejacent is true in some or every possibility compatible with the contextually-relevant information state. The range of possibly contextually relevant information states that determine the set of possibilities quantified over is pretty big. At one end of the spectrum lies the solipsistic reading, the reading where the BEM is used to report solely on the speaker's information state. Here is an old example from Kratzer (1986):

[4] This is the now canonical analysis of the logical structure of modal statements developed by Kratzer (1977, 1981, 1991).

[5] As in most of the literature, we will not attempt to spell out a compositional semantics of how such phrases get to determine the value of *B*. And, when this won't be too confusing and when style demands it, we'll also skate over the difference between *B* and its denotation (at a context–index pair) $[\![B]\!]^{c,i}$.

> Suppose a man is approaching both of us. You are standing over there. I am further away. I can only see the bare outlines of the man. In view of *my* evidence, the person approaching may be Fred. You know better. In view of *your* evidence, it cannot possibly be Fred, it must be Martin. If this is so, *my* utterance of (4) and *your* utterance of (5) are both true.
>
> (4) The person approaching might be Fred.
> (5) The person approaching cannot be Fred.
>
> Had *I* uttered (5) and *you* (4), both our utterances would have been false.

Kratzer claims that the BEMs in this example are claims about the speaker's evidence. When we consider that the first speaker may well be aware that the second speaker has a better vantage point and may thus have a better idea about who is approaching, it makes sense to assume that the first speaker is only making a claim about *her own* somewhat limited information state.

But the speaker-centric interpretation is not the only one available. If it were—as DeRose (1991) points out—it would never make any sense to say *I don't know whether it might be that p*. But it can make sense to say that. John has had a screening test that can rule out cancer but will not determine that he has it if he does. After the test has been run and the doctors have the results, Jane can say things like

> (6) I don't know whether John might have cancer; only the doctors know. I'll find that out tomorrow when the results of the test are revealed.

And if the speaker-centric interpretation were the only one available, we would expect to be able to gloss (6) with a sentence in which the restricting *in view of* phrase picks out Jane's knowledge. But that isn't so:

> (7) ?? I don't know whether in view of what I know John might have cancer; only the doctors know. I'll find that out tomorrow when the results of the test are revealed.

Whatever information state this *might* quantifies over, it doesn't include just Jane's knowledge.[6]

More objective readings—readings in which the modal quantifies over a modal base that goes beyond the speaker's information state—need to be available. One way of achieving greater objectivity is by allowing modal bases to provide possibilities compatible with the information that a *group* of agents has. For example, it is plausible that Jane fully intends her *might* to be, in part, about the information John's doctors have. But we may need more. Sometimes a *might*-claim seems false even if, in the context, the prejacent isn't ruled out by anything the speaker or hearer or any pooling thereof knows—for instance, if that information was there and they made some mistake in

[6] Another case in which we clearly don't have a purely solipsistic reading can be manufactured out of Kratzer's scenario above. The speaker who is further away might well ask the plausibly better informed person *Might the person approaching be Fred?*, which could not be exclusively about her own information state.

not seeing it and thus *should* have known better.[7] Holmes never botches an investigation. But his lesser-known cousin Schmolmes sometimes does. Reading through his interview notes, Schmolmes makes some errors in his deductions and declares

(8) Ah, the gardener might be the culprit.

Alas, the gardener not only didn't do it, but—as Schmolmes's own interview notes conclusively establish—he couldn't have. Poor Schmolmes just got confused and didn't connect all the dots. Even if his conversational partners aren't any the wiser, and so even if there is no plausible way of drawing the contextually relevant group boundaries to include anyone who knows that the gardener didn't do it, we still get the impression that Scholmes said something false.[8]

The canon as such is non-committal about how context can decide the relevant body of information that BEMs quantify over. So it is compatible with insisting that the relevant body of information represents some pooling of the information available to the relevant group of agents and it is compatible with insisting that the relevant body of information represents not merely what those agents know but what they may come to know—what is, borrowing Egan's (2007) nice phrasing, within their epistemic reach.

The basic lesson is that these amendments are not really deviations from the canon: it has plenty of room for this kind of flexibility in readings available for BEMs. The canon is thus a big tent, with lots of room for family squabbles about just how context decides to pool information and just what counts as within an agent's epistemic reach. Since our main point here is independent of how those disputes get resolved, we don't want to take sides and so we consider the canon in its most general flavor.[9]

[7] The point was first made by Hacking (1967) and then by Teller (1972) and then by DeRose (1991).

[8] Well, actually, the judgments here aren't uniform. As in other spots, this variability in judgments is itself something that needs to be explained by our best theory. So it is a mistake to tailor the theory—canon or otherwise—to automatically deliver the verdict that (8) is false.

[9] We do want to register one thought, though. It is hard to cash out what, in a context, counts as "within epistemic reach." Hacking says it's "practicable investigations" that count; DeRose says it's "contextually relevant" ones that do. But neither of these ways of constraining the notion seems right. For it is easy to find cases in which there is a "practicable" way of finding out the truth of the matter about φ and yet *might* φ and *might not-*φ both seem true. An example from Teller (1972):

(i) It might be a girl and it might be a boy. Should I buy blue or should I buy pink?

Doting soon-to-be grandmas say things like this—and thereby speak truly—even though there is a practicable test that can give them just the information they lack to help inform their buying behavior.

And saying that it is "contextually relevant" ways of coming to know seems no better. For it is easy to find cases in which there are relevant ways of coming to know that φ is false and yet *might* φ seems true. Alex is helping Billy in the search for her lost keys and says:

(ii) They might be in the car.

The keys, in fact, are not in the car. Has Alex thereby said something false? Not to our ears. Does that mean that checking the car does not count, in this context, as a "relevant way" of coming to know? No, since the point of Alex saying what she did is precisely to get Billy to check it out. Consulting Schmolmes's interview notes can count, as can ships logs (Hacking), and medical test results in a sealed envelope (DeRose). But performing a baby gender test does not, and apparently, neither does looking in the car. The project of gerrymandering epistemic reach to fit these boundaries has all the hallmarks of a project we wouldn't want to take up. So we would instead just take BEMs to quantify over the information held by a contextually relevant group of agents, and be quite egalitarian about the kinds of agents we

But we do want to insist that the pooling of information be constrained in sensible ways, and that does throw some versions of the canon out of the tent. The only context-dependence at issue for us here is the dependence of BEMs on the possibilities compatible with the information a relevant group has. So, where c is a context, we will often write G_c as the c-relevant group. And when it suits our purposes we will sometimes talk as if contexts just are the sets of agents themselves, eliminating c altogether in favor of G.

We might—as Hacking and DeRose both do—say that a BEM requires that no member of G_c know that the prejacent is false. Putting things the other way around: it is compatible with what each knows that φ. Let's let f_x be a function from indices i (worlds) to the set of indices compatible with what x knows at i. (Thus we will assume that f_x is both reflexive and euclidean.) Formally, this proposal is:

(9) $[\![\textit{might}\ \varphi]\!]^{c,i} = 1$ iff $\forall x \in G_c : \exists w \in f_x(i)$ such that $[\![\varphi]\!]^{c,w} = 1$

But now we have lost the idea that a modal is a quantifier over a modal base: there is no one set of possibilities throughout which we check for some φ-worlds. That is too high a cost to pay, for it makes epistemic modals different in structure from the other modals in the language. So that is ruled out. Better to aggregate in some other way, pooling the information states of the members of G into some aggregated information state, and take BEMs to quantify over possibilities compatible with that. That would be to figure out a denotation for the modal base B in terms of what it is each $x \in G$ knows. In that case, we can simply stick with the initial modal-base skeleton that says that *might*, in a context and at an index, is an existential quantifier over the modal base $[\![B]\!]^{c,i}$. All we need is to find a good candidate for B in terms of G.

There is a spectrum of group-level modal bases—candidate values for B—that can be built from individual selection functions representing what the members of a group know.[10] But since we do not want to get involved in the family disputes about which is the right one, we will instead just insist on the following constraint:

(10) AGGREGATION:
Suppose c determines B by determining G and c' determines B' by determining G'. Then $G \subseteq G'$ implies $[\![B']\!]^{c',i} \subseteq [\![B]\!]^{c,i}$.

The more inclusive a group, the harder it is for a *might*-claim that quantifies over what it knows to be true. As an example of a story that implies (10): suppose we ignore the issue of epistemic reach (opting instead to be egalitarian about what kinds of things get to count as agents in a context), and suppose we pool information available to those in G by distributing it:

recognize: stores of information—ships logs, interview notes, and computers—can *ceteris paribus* count as "agents", but plain facts in the world—that the baby will be a girl, that it isn't raining, and that the keys aren't in the car—cannot. Exploring this feature of epistemic modals—the resistance to plain facts but sensitivity to stores of information—is a task for another occasion. We have some preliminary remarks about this feature, which makes epistemic modals related to expressions of evidentiality, in von Fintel and Gillies (2007).

[10] The classic reference for multi-agent epistemic logic is Fagin et al. (1995).

(11) Fix a c-relevant group G_c. Then $[\![B]\!]^{c,i} = \bigcap_{x \in G} f_x(i)$

Then (10) follows straightaway.[11]

So the canon is quite flexible. It captures solipsistic readings—which are almost always available—as a special case. If $G = \{\text{speaker}\}$ then what is compatible with what that group knows is just what is compatible with what she knows. And there is nothing in any of this that requires a hard-wired commitment that the speaker is always a member of the relevant group provided by context. To put it provocatively: just because s is in a context c (she's the speaker in the context of utterance after all) it does not follow that s is at all *relevant* in c and so for all the canon has said she may not be in G_c. Similarly, there's nothing that requires the members of G_c to be parties to the conversation s is having.

All of this is in the canon's favor: given the flexibility of what gets to counts as a relevant group in a context, that makes for a pretty broad spectrum of information states that can be relevant and thus for predicting flexibility at just the spots that the data demand it. It is thus tempting to think that BEMs quantify over the pooled information in a group and that the context determines what the relevant group is. So far so good. Since we like to give in to temptation, that's what we want to think. It's just that we have found a problem with the second conjunct: that *the* context *determines* what the relevant group is.

3. A Realistic Scenario

Assume that the canon is right: BEMs quantify over possibilities compatible with the pooled information of some relevant group. Given the broad range of possible meanings—all the way from solipsistic readings to ever more encompassing group readings—we would then expect that contexts are made to work hard to resolve this indeterminacy. But a look at realistic scenarios makes us suspect that contexts in fact do not bother much with resolution at all.

Alex is aiding Billy in the search for her keys:

(12) Alex: You might have left them in the car.

[11] Distributed knowledge is also a rather intuitive notion. Here's the comment of a 9/11 widow (heard on NPR on 4/11/2004) when the 9/11 commission report was discussed:

(i) We knew more than is being owned up to. But nobody put the pieces together.

There are other ways of "pooling" information that can plausibly lay claim to being distributed information. Here is a natural one: the information every member of the group would have after all of the members (successfully) share what they know. This notion coincides with (11) for non-modal prejacents, but diverges thereafter. Example: we know that p and that you do not know it. That is a bit of information that is distributed in our group in the sense of (11), but not something any of us will know after we all share what we know. For then you *will* know that p (and we'll know that you do). There are, of course, other group-level knowledge operators—what is common knowledge among G, what everyone in G knows, what someone in G knows, and so on—but none that are clearly as plausible as what is distributed among G for the purposes of BEMs.

From here the conversation can take one of two paths. If Billy cannot rule out the possibility raised by Alex, an appropriate response might be:

(13) Billy: You're right. Let me check.

On the other hand, if Billy *can* rule out the prejacent, we find responses such as:

(14) Billy: No, I still had them when we came into the house.[12]

This is a perfectly ordinary scenario between speaker and hearer, and there are two natural ways the dialogue could go. What reading does the BEM that they are discussing have? What resolution of the modal base does the context of their conversation supply?

For simplicity, let's consider just two candidate resolutions of the contextual parameter. Under one resolution, the relevant group has as its only member the speaker Alex. Under this speaker-centric resolution, $G = \{\text{Alex}\}$. Call the associated reading the *A*-reading. With respect to this disambiguation, the BEM is about Alex's evidence. The other resolution we'll consider takes both Alex and Billy to be relevant. Under this groupwise resolution $G' = \{\text{Alex}, \text{Billy}\}$. Call the associated reading the $A+B$-reading. Here the BEM is about the information that the non-trivial group of Alex and Billy has. So, which of these two candidate resolutions is the one that the context of their conversation determines to be the one giving the relevant interpretation of Alex's BEM ? We will now see that in fact neither reading can be the one that is at issue here.[13]

The speaker-centric resolution faces trouble right out of the gate. When Alex issues her BEM, both potential replies by Billy are open, one registering denial or disagreement and the other registering acceptance or uptake. But, assuming the speaker-centric resolution and thereby assuming that the dialogues traffic in the *A*-reading, what is it that Billy is denying or taking up?

First, take (14). What could Billy be disagreeing with by uttering this? Not the content of the *A*-reading: Billy has no reason to doubt that Alex's information did not rule out that the keys were in the car. Rather the "No" of Billy's answer is—plausibly—a negation of the prejacent: the keys are not in the car. Since the truth of the prejacent is ultimately what matters here, it is unsurprising that Billy would deny it directly.[14] But the reply in (13), *You're right*, is also open. What could Billy be taking up here? Again, not the content of the *A*-reading: Billy is in no position to comment on whether Alex's information state leaves it open that the keys are in the car. But

[12] Notice that it is virtually obligatory for Billy to justify the rejection of the BEM by giving the key bit of evidence that leads him to rule out the prejacent. We actually don't know exactly how to derive this obligation, even if it seems blindingly obvious.

[13] We do not even consider the possibility here that the BEM is meant under the *B*-reading. It is obvious that *A* has no business making assertions about *B*'s information state. Later, we will resurrect the *B*-reading, however.

[14] The possibility that reactions to a BEM target not the modal claim but the prejacent is one that always needs to be kept in mind when using dialogues as data for the semantics of epistemic modals. At least some of the recent relativist literature is careless in this regard.

unlike the denial, targeting the prejacent here is not a plausible dodge. Whatever Billy is agreeing to, it is not that the keys *are* in the car. In fact what it *feels* like is that Billy is agreeing that it is compatible with what they *qua* group know that the keys are in the car. But that is the $A+B$-reading, not the A-reading.

And though Billy's space of responses is quite open, it is not unconstrained. One thing she definitely can't do is reply

(15) Billy: # OK, but *I* know that they're not there.

If the dialogues here trafficked in the A-reading—if the context resolved our group down to just the speaker Alex—then Billy's response in (15) should be just fine.

That would seem to leave us with the groupwise resolution of the contextual parameter and the associated $A+B$-reading. This resolution would seem to be compatible with both replies. If Billy doesn't know where the keys are, then her reply in (13) would be on the mark since in this case there would be keys-in-the-car possibilities among those compatible with what G' knows.[15] And if Billy knows they're not in the car, then since she is a member of G' her private knowledge is sufficient to guarantee that there are no such keys-in-the-car possibilities among the possibilities compatible with what G' knows. The denial in (14) is thus also on the mark.

But we've hopped out of the frying pan and into the fire. For under this resolution, it is hard to see how Alex could be in a position to assert (12) in the first place. She does not seem to be within her linguistic rights to be claiming that the group's information cannot rule out the prejacent.[16] After all, Alex does not know whether Billy has private information about the whereabouts of the keys. So if the usual norms of assertion apply in the case of epistemic modals, then—no matter whether your favorite[17] story for assertion requires justified belief or knowledge or truth of the asserted content—Alex has no business asserting the $A+B$-reading, unless she is confident that Billy has no information that goes beyond her own information. That does not seem right: in our scenario, Alex can utter the BEM without such confidence in place. (And if she did have that confidence, we wouldn't need to resort to the $A+B$-reading at all.)

Making sense of Billy's space of available replies seems to close off resolving the relevant group to just Alex and pushes us to the group reading. But making sense of Alex's assertion goes just the other way, closing off group readings as being beyond what she could reasonably assert (if the usual norms are in force at any rate). So neither

[15] Once we look a bit closer at this, we'll see that what we say here isn't quite true. Just because neither Alex nor Billy can rule out that the keys are in the car doesn't mean that the group of the two of them can't rule that out. We will explain later why Billy can still agree to the $A+B$-reading.

[16] As far as we know, the earlier literature on group readings of BEMS (DeRose 1991; Teller 1972) did not discuss the question of what justifies the assertion of such a reading by a speaker who does not have unusual access to the other group members' information state. Perhaps, it is not an accident that DeRose's paradigm case of an indisputable group reading is an embedded (unasserted) BEM: *I don't know whether John might have cancer*.

[17] We don't pick sides in the fight; see Williamson (1996), Weiner (2005), and Lackey (2007) for some of the contestants.

of the predicted readings is the one that the context determines as the interpretation of the BEM that Alex uttered.[18]

This is where any CIA agents who are eavesdropping on our discussion are hopping up and down impatiently. They advise that this is precisely the reason why their analyses are to be preferred: the contextualist analysis cannot make sense of the apparent observation that as far as the speaker is concerned, *the* context provides the solipsistic resolution of the context-dependency, while when the hearer comes in, *the* (*very same*) context supplies the group reading. So, why not say, they insinuate, that the way the modal is read varies with the context of *assessment*, which varies with who the assessor is: the speaker in the moment of utterance, the hearer (or the group) in the moment of reception? Well, we've already explained why not (in von Fintel and Gillies 2008). Here, we will show how the contextualist analysis can respond to the puzzle. The story we will tell will take the semantics of BEMs to be pretty much as the canon says: they are quantifiers over possibilities compatible with the information a contextually relevant group has. But we will tell a much richer story about the pragmatics of BEMs.

So, what does the BEM really mean here? Our contention is that rather than having *neither* of the solipsistic or group readings (and thus having some as of yet novel reading, perhaps such as the one promoted by the CIA agents), Alex's BEM actually has *both* readings—possibly many more, in fact—and that this kind of multiplicity of meanings is precisely what gives BEMs their peculiar properties. The context does not, in general, determine what the relevant group is. Instead, it leaves this underspecification intact, and—we will say—epistemic modals exploit this.

4. Ambiguity by Design

The canon requires contexts to do more than they in fact do: a context in which a BEM is deployed need not, it seems, fully determine a relevant group of agents. This is more feature than bug, though, generating ambiguity as if by design.[19]

Not all actual contexts of use for a natural language sentence will successfully resolve all indeterminacies and ambiguities of that sentence. That is well known. Sometimes, of course, this will cause the conversation to break down and will thus necessitate a negotiation between the participants over the interpretation of the sentence. When Bond and Leiter are looking at a group of members of parliament and Leiter suddenly shouts:

[18] A possibility we will not pursue as such here is that the contextually determined content of the BEM is in fact the $A+B$-reading, but that it is put forward by Alex not as an assertion but as a weaker kind of speech act. In our story, that is almost right: Alex does weakly put the $A+B$-reading in play but what she is in a position to assert is the A-reading. Stay tuned, we'll explain soon.

[19] Our account was partially inspired by a remark by Angelika Kratzer (p.c. to Kai von Fintel, at a UMass linguistic colloquium in December 2003). Roger Schwarzschild's talk (1999) on cases where contexts underdetermines context-dependent items was also important to us. Neither should be held responsible for what we did with their ideas.

(16) That guy is an assassin. Shoot him before he can do anything.

one presumably expects Bond to ask for a more precise reference as to which of the politicians he's supposed to take out. But other times, we proceed without full resolution of contextual ambiguities. Bond is stealthing his way through a dark corridor, his local guide in tow. A figure moves in the shadows and the guide shouts:

(17) Watch out! He's trying to shoot you.

Bond does not puzzle over the referent of that hanging anaphor; he ducks (and, amazingly, disarms the figure in one motion).[20] And yet other times, the indeterminacy is there as if by design, and this is what we claim is going on with BEMs.[21]

One more example. Billy meets Alex at a conference, and asks her:

(18) Where are you from?

That question is supposed, given a context, to partition answer-space according to how low-level in that context Billy wants his details about Alex to be. But notice that it's not really clear whether Billy wants to know where Alex is currently on sabbatical or where Alex teaches or where Alex went to graduate school or where Alex grew up. And—the point for us—Billy might not know what he wants to know. He just wants to know a bit more about Alex and will decide after she answers whether he got an answer to his question or not. He doesn't have to have the level of granularity sorted out before he asks the question. So context (or context plus Billy's intentions) need not resolve the contextual ambiguity.[22]

When a BEM is deployed, the facts about the conversation up to that point might be compatible with multiple ways of drawing the boundaries to what can plausibly count as "the relevant group." Since we are pretending that this is the only relevant contextually supplied information, that means we can think of utterances taking place against a cloud of admissible contexts—one for each resolution of the relevant group that is compatible with the facts as they are when the BEM is issued. And indeed, we will assume that there is a one–one correspondence between admissible contexts and potential resolutions of the relevant group.

It is important to realize that the proposal is not that some kind of objective context does provide a determinate resolution of the BEM and that the conversational parties are ignorant of or indifferent towards what the context is. There is no such thing as "the context", only the contexts admissible or compatible with the facts as they are. The context of the conversation really does not provide a determinate resolution and

[20] The example in (17) is due to van Deemter (1998). Idealizing, we might say that Bond is employing the strategy of diagonalization here (Stalnaker 1978): he computes the proposition that whoever the guide is referring to is trying to shoot him.

[21] It has been observed that sometimes ambiguous sentences are used purposefully to convey more than one proposition, namely in the context of jokes (Raskin 1985), poetry (Su 1994), and other less than straightforward uses of language; see Poesio (1996) for discussion. Our proposal here is that a multiplicity of meaning is also detectable in at least ostensibly more straightforward uses of language.

[22] We owe this example to Chris Potts (p.c.).

we propose to model this by saying that there is a cloud of contexts at the given point of the conversation.

There are different ways of making this intuition more precise, each way representing a different way of distributing the labor between the semantics and the pragmatics. One way is to allow the possibility that some constructions—in our case, BEMs—have as their semantic value not propositions but *sets* of propositions. And to do that we could insist that semantic values are assigned not with respect to contexts but with respect to sets of admissible contexts. With a little care this can be done in such a way that we do not have to seriously complicate the pragmatics.

But there are other options. The one we opt for here makes for a better view of some of the landmark properties of BEMs. The point for us is that there is a space of pragmatic stories that can be told that say how the contextual ambiguity of BEMs gets exploited in a way that does justice to the data about them. We tell one such story, giving some basic pragmatic principles that combine with the basic semantics for BEMs to do that work.

We begin by saying what *travels* in a conversation, or what proposition(s) the speaker of a BEM *puts into play*. When a proposition is put into play, it's available for denial and uptake. When an utterance is contextually underspecified, the propositions corresponding to the various disambiguations are put into play.

(19) TRAVEL:
Suppose the facts (linguistic and otherwise) up to t allow the groups $G_1, G_2, \ldots$ as resolutions of the contextual parameter, these resolutions delimiting the cloud C of contexts. Then an utterance of $might(B)(\varphi)$ with respect to C at t puts into play the set of propositions P such that for some $c \in C$: $[\![might(B)(\varphi)]\!]^c = P$.

This means that, as far as the semantics is concerned, there is no underspecification here: BEMs get assigned normal semantic values at (determinate) contexts. But since utterances of them take place against a cloud of such determinate contexts—since there is a set of ways compatible with the context of determining a contextual parameter—those utterances put in play a set of such semantic values. This would make our proposal here a relative of—how close we shall not guess—proposals for how the grammar deals with other types of underspecification.[23] Since multiple propositions travel following a single utterance of a BEM, there is a lot of explanatory work left for the pragmatics to do. We will look at the pragmatics of this ambiguity by design from both sides of a conversation. We will again use the simple but realistic dialogue between Alex and Billy about the whereabouts of the keys as our test case.

When Alex says *The keys might be in the car*, this BEM puts into play multiple propositions, each of which makes a claim about the information state of a group engaged in the investigation. There are arguably three such groups: the singleton groups {Alex},

[23] Poesio (1996), for example, gives a semantics for an underspecified language that assigns sets of standard meanings to expressions of that language.

{Billy}, and the non-trivial group {Alex, Billy}. So, there are three readings put in play: the *A*-, *B*-, and *A*+*B*-readings. Now, what are speaker and hearer doing with that set of propositions?

First, let's consider the speaker. She manages to put in play a set of propositions. But she does not have to be in a position to assert each one. When Alex uttered *The keys might be in the car* with its three meanings, she was not in the right position to flat out assert either the hearer-centric *B*-reading or the groupwise *A*+*B*-reading. Since those propositions are among the set of propositions she put in play, and since her utterance was appropriate, that means that she didn't have to be in a position to flat out assert each of the meanings the BEM in her mouth had.

But perhaps this relies on a mistaken theory of what the norm of assertion comes to. Perhaps, according to the right story, it turns out that when a speaker utters an underspecified sentence she must—deontic *must*, that is—be in a position to assert each of the propositions she puts in play. And perhaps someone would enjoy arguing for all that. But not us. We'll leave the norm of assertion untouched and we will say that the speaker only has to be in a position to flat out assert *one* of the propositions she puts in play; any one of them will do:

(20) ASSERT:
Suppose an utterance of *might*(*B*)(*φ*) by *S* puts in play the propositions $P_1, P_2, \ldots$. Then *S* must have been in a position to flat out assert one of the P_i's.

Our proposal is that in order for a speaker to be within her linguistic and epistemic rights when she issues a BEM against a cloud of contexts, she has to be in a position to flat out assert one of the meanings it can have, given that cloud.

Return to our little dialogue. When Alex utters the BEM, with its three meanings, there needn't be a fact of the matter as to which of the three meanings she intends to assert. But she needs to be in a position to flat out assert at least one: the *A*-reading, the *B*-reading, or the *A*+*B*-reading. Given the facts of the scenario, the minimal requirement is that Alex needs to be in a position to assert the *A*-reading. That will be the weakest reading she will stand in the appropriate relation to. In other words, given the facts of that scenario, Alex is justified in uttering the BEM iff she is justified in claiming that her evidence does not rule out the prejacent. As far as the norms of assertion go, it's as if she had uttered an explicit claim about her own evidence. But that's not what the BEM in her mouth means: it has the three meanings at once.

Things can be different in different scenarios. Part of what's right about the canon—and part of what's wrong with the CIA—is that it can be perfectly sensible to assert *might φ* even when you know that *φ* is false. Pascal and Mordecai are (still) playing Mastermind.[24] After some rounds where Mordecai gives Pascal hints about the solution, Pascal asks whether they might be two reds. Mordecai answers:

[24] We've been using this example, to make various points about epistemic modals, quite a lot lately: see e.g. von Fintel and Gillies (2007, 2008, 2010).

(21) That's right. There might be.

He can answer this way even if he knows there aren't two reds. As far as the norms of assertion go, it's as if he had uttered an explicit claim about Pascal's evidence.

Now, let's consider the hearer's side of the exchange. Since a BEM puts multiple propositions in play, we need to sort out which of these a hearer should react to, which a hearer takes as the appropriate target for uptake or denial. Again, we say that in order to take up and accept a BEM issued just prior, a hearer must be in the right relation to just one of those propositions. But now not just any one of them will do. Instead, we argue that the hearer is guided by what response to which proposition will be most informative in the conversation. When the modal is an existential like *might*, this will in fact lead to a dominance of negative replies.

(22) CONFIRM/DENY
Suppose an utterance of $might(B)(\varphi)$ by S puts in play the propositions $P_1, P_2, \ldots$. Then a hearer H can confirm (deny) the BEM if the strongest P_i that H reasonably has an opinion about is such that H thinks it is true (false).

In our dialogue, the BEM that Alex utters has the A-reading, the B-reading, and the $A+B$-reading. Alex is justified in uttering the BEM because she is in a position to flat out assert the A-reading. But just asserting the A-reading isn't what she's doing. When Alex puts the three propositions in play, the other readings, the hearer-centric reading and the strong group reading, are floated. She does not have to be in a position to assert those in order for them to be available for Billy to react to. It is as if she is conjecturing that the B-reading and the $A+B$-reading are true or asking whether they are true. Billy confirms/rejects the BEM based on those readings. Billy reasonably has an opinion on the B-reading, and especially if he can rule out the prejacent, he also reasonably has an opinion on the $A+B$-reading, namely that it is false (if he can rule out the prejacent, then so can any group to which he belongs).

One might have thought that even if Billy can rule out the prejacent based on his private information, a principle of charity should lead him to accept the BEM under the A-reading. After all, there is ambiguity one resolution of which would leave Alex having spoken truly. Instead, our principle has it that the more cooperative thing to do is to reject the BEM because it is false under the $A+B$-reading. At least, this is *ceteris paribus* the right thing to do—for instance if the goal of the conversation is to ultimately determine the truth of the prejacent or even the answer to a more general question. In our case of the misplaced keys, the ultimate goal is to find out where the keys are and the proximate goal is to figure out whether they are in the car. It is not the goal of the conversation in any important sense to find out whether the speaker's evidence or the group's evidence at the time of the conversation rules out that the keys are in the car. So, the proper thing to do—the more cooperative conversational

move—is to deny the BEM under the $A+B$-reading and by entailment thus deny the prejacent.[25]

Let us look a bit closer at the case where Billy cannot rule out the prejacent on the basis of his information. Obviously, he will confirm the BEM (*You're right*). Which reading of the BEM is he reacting to? Our principle says that it is the strongest reading he reasonably has an opinion about. Which one is that? It is obvious that he reasonably has an opinion about the B-reading. What about the stronger $A+B$-reading? Can he reasonably have an opinion about that?

We have said that the group readings of BEMs make claims about the pooled knowledge of the relevant group. When Alex utters the BEM, Billy concludes that Alex is not in a position to rule out the prejacent. If Billy himself is also not able to rule out the prejacent, he might put 2 and 2 together and conclude that the group of the two of them cannot rule out the prejacent. But wait: while 2 and 2 makes 4, the group reading may actually be more like 4.5: suppose both Alex and Billy know that q and r entail $\neg p$, but just Alex knows that q and just Billy knows that r. Then p is compatible with what each knows, but not with what they know *qua* group if the pooling goes by distributed knowledge as in (26). Now even if Billy learns that p is compatible with what Alex knows, he still can't rule out this sort of thing. And so, after learning that the A-reading of a BEM $might(B)(p)$ is true, Billy does not know that the $A+B$-reading is true.

Nevertheless, we think that in many cases where B cannot rule out the prejacent, it is in fact reasonable for B to jump to the conclusion that the group reading is true—even if strictly speaking that group reading does not follow from that. That is, it tends to be a mutual expectation in a conversation that partners in it are similarly situated. That expectation is both rough and defeasible, but still ripe for being appealed to in a conversation:

(23) DEFEASIBLE CLOSURE
If H knows that φ is compatible with what x knows, for each $x \in G$, then it is reasonable for H to defeasibly infer that φ is compatible with what G knows.

This is a merely defeasible inference, since there are certainly cases where one shouldn't draw the inference: cases where the stakes are high, cases where one has reason to think that others in the group are far more informed about the prejacent than oneself and those from whom one has heard, and so on. We hypothesize that this kind of presumption is suitably conventionalized, that conversational partners naturally expect each other to draw on it, and that when the conditions are not ripe for it they expect this fact to be reasonably transparent to each other.

[25] Instead of justifying our principle in (22) by telling a story about cooperativity, we could also simply stipulate that a hearer confronted with a systematically ambiguous sentence whose meanings can be ordered in strength should treat the utterance as if it carried the strongest of those meanings. This would align our proposal with other work that has argued for a STRONGEST MEANING HYPOTHESIS (Dalrymple et al. 2008; Winter 2001).

If we are in a context where it is reasonable to draw the defeasible closure inference and if *B* cannot himself rule out the prejacent, then *B* upon hearing *A*'s BEM can reasonably have an opinion on the *A*+*B*-reading, namely that it is true. Hence, *B* can confirm the BEM on that strong reading.

Now, let us zoom out to a bird's-eye view of the BEM exchange: the BEM itself has three meanings, but Alex acts as if the solipsistic reading were the one that matters, while the hearer acts as if the hearer-centric or even the group reading were the active ones. This asymmetry is what gives BEMs their quasi-magical properties: a speaker can utter them based on just her own evidence but it serves as a probe or test or trial balloon into the hearer's evidence. When things go well and a hearer takes up a BEM, this fact becomes common ground between speaker and hearer and thus it follows that it is common belief between them that the prejacent is compatible with the information that they *qua* group have.[26]

It is easy to see why Alex might want to float those (possibly) stronger readings. If what is important is where the keys are, and not what's compatible with what who knows, then finding out the truth about those stronger readings is a pretty clear path to finding her keys. But none of that is to say that the BEM in the speaker's mouth was an assertion of that strong group reading. Take our toy example once more. If what matters goes beyond where the keys are, further moves are imaginable. If Alex and Billy are bickering already, one could encounter a dialogue like this one:

(24) A: The keys might be in the car.
B: They're not. I still had them when we came into the house. Why did you say that?
A: Look, I didn't say they *were* in the car. I said they *might be* there—and they might have been. Sheesh.

Here, Alex is sticking to her guns, defending her BEM on the basis of a weaker reading than the *A*+*B*-reading. Once she does this, there is no basis for a continued dispute and the only avenue open to Billy at this point is to back off.

Our story about the ambiguity-by-design of BEMs would perhaps be even more convincing if we had independent confirmation that such a story is useful beyond epistemic modals. An obvious place to look are contextual restrictions on other quantifiers:

(25) Every student was at the meeting.

It does not seem implausible to say that the precise delimitation of the contextual domain of quantification for (25) can often be indeterminate in a realistic context. We suggest that the speaker of (25) in such a context has put into play a set of propositions

[26] Note that because of the strength of the distributed knowledge reading of group-BEMs, it is usually not common *knowledge* that the group can't rule out the prejacent. This is why we are saying in the text that after the BEM exchange is completed, it is merely common belief (or even only common acceptance) that the group can't rule out the prejacent. But that's OK since a group can be wrong in thinking that a BEM is true.

that differ in the domain of quantification. While we can't elaborate on this suggestion here, we should point out that in such cases, it seems more often appropriate for the speaker to retreat to a weaker resolution of the domain than it does in the case of BEMs:

(26) Sally: Every student was at the meeting.
George: What, even those that are on leave in Nicaragua?
Sally: No, what I meant was every student in residence.

In the case of BEMs, defending the BEM under the weakest, solipsistic interpretation is often—though not quite always—irrelevant (exceptions include: cases of antagonistic conversations, being on the witness stand, and playing Mastermind).

Let us end the presentation of our basic story by thinking about the conversational dynamics one more time. It might seem outlandish to claim that BEMs are typically multiply ambiguous and that this multiplicity is dealt with not as a defect but is embraced and processed quietly and without fuss. But we hope to have shown that the pragmatics of such exchanges relies on some rather intuitive principles. It should also be noted that once a hearer has confirmed or denied the BEM with all its indeterminacy, the resulting common ground is quite determinate. If *B* denies the BEM—and, of course, assuming the conversation does not there derail—then what's common ground is that the prejacent can be ruled out by the group, and hence that the prejacent is false. If, on the other hand, *B* confirms the BEM, then there are two possibilities, depending on whether the conditions are right for the principle of DEFEASIBLE CLOSURE to apply. If the standards of the conversation are too strict or there are other defeaters nearby, then there is no appeal to DEFEASIBLE CLOSURE: instead it is common ground that no member of the group can rule out the prejacent. Hence what becomes common ground is, in effect, the kind of reading of the BEM that Hacking and DeRose posit *even though that is not available as something 'might' might mean*. And if DEFEASIBLE CLOSURE *does* apply, then what is common ground between them is the distributed group reading: that their information *qua* group cannot rule out the prejacent. Again, what becomes common ground is not something any of the group members flat out asserted, but is something much stronger. That is valuable, and what exploiting ambiguity by design can buy us. There is no lingering uncertainty one would have to worry about. We hope that it has become clear that BEMs with their multiplicity of meanings are a very useful device to have in one's grammar.

5. Beyond Alex and Billy

We have been looking at a very simple but still realistic scenario involving a single speaker and a single hearer trying to ascertain the whereabouts of some keys. Somewhat realistic, but still an idealization. So we shall take our literary license and give Alex and Billy another roommate, Chuck.

All three roommates are looking for the keys. Alex again says that the keys might be in the car. In this context, possible resolutions for the BEM are the solipsistic readings (the *A*-,*B*-, and *C*-readings) plus the relevant group readings (the $A+B$-, $A+C$-, and $A+B+C$-readings).[27]

As before, Alex can utter the BEM because she is in a position to assert the *A*-reading. But at the same time, she is floating the other readings. And now, we consider the possible responses from Billy and Chuck. Our CONFIRM/DENY principle (22) says that each of them should determine their reaction to the BEM based on the strongest relevant reading about which they can reasonably have an opinion. If either Billy's or Chuck's information state rules out the prejacent, then we expect the same denial as before. And assuming the conversation does not derail, or that there is no conversational retreat like we saw in (24), then it becomes common ground between them that the keys are not in the car.

But suppose neither Billy nor Chuck has information that eliminates the possibility that the keys are in the car. We will assume that the conditions are ripe for applying DEFEASIBLE CLOSURE. (If they aren't, they confirm the BEM, and the fact that the prejacent is compatible with what each knows is common ground.) Both will then confirm the BEM. That is because each has good reason for the strongest available meaning about which he has an opinion. Billy because he believes the $A+B$-reading, and Chuck because he believes the $A+C$-reading. After both confirm the BEM (assenting, in effect, to different floated meanings of it), then it is common ground that the prejacent is compatible with Alex's information, Billy's information, and Chuck's information. All parties can then appeal to DEFEASIBLE CLOSURE once more to infer, defeasibly, the $A+B+C$-reading. And since they can expect the others to do the same, this can become common ground.

But this two-step procedure isn't always required. Sometimes conversational partners can jump straightaway to confirming the strongest group reading floated. A team of investigators has been collecting clues at the scene of the crime, when the Detective calls them all together:

(27) a. Detective: Do we think the murderer might have used an icepick and slipped out the window?
b. Investigator #1: Yes, he might have.

The investigator looks like she is answering Detective's question about what is compatible with what the whole team of investigators has been able to find out so far (he asked what *we think* about the possibility). That means she is not confirming the Detective-plus-Investigator #1-reading, but a much stronger one. And if none of her co-investigators complain, she will have done so felicitously (though she could still turn out to be mistaken). How did she do that?

[27] Given the set-up, it seems unlikely that Alex's utterance puts in play a $B+C$-reading. But in other scenarios it might: for example, if Alex knows full well where the keys are but her job is to test Billy and Chuck's key-sleuthing skills Billy can ask *Might the keys be in the car?* and Alex can truthfully reply *They might be*.

Answering collective queries or confirming group conjectures can be easier than you might have thought. We are at Hullabalooza when Rock Star takes the stage and asks:

(28) Rock Star: Is Springfield ready to rock?!?

The two of us are, in fact, ready to rock. But we don't know about everyone else. Plus, we are not the town spokespeople, and neither is any one person in the audience. What Rock Star wants to know is if the group is—*qua* group—ready to rock and that is something none of us could answer on our own. And he knows that. So why would he ask? Because we can *anticipate* that each of us is ready to rock, and then answer on that basis that the group is. It is excellent evidence that Springfield is ready to rock if, when he asks, we all yell *Yes!*. Excellent, but not quite conclusive: that is because, plausibly, there are group-level facts about our readiness to rock. We can all be individually ready, but not gel in the right sort of way *qua* group.

Of course, Rock Star has a hunch. He suspects that his query will be met with a raucous chorus of *Yes!* from the audience. It seems acceptable for each of us to confirm Rock Star's suspicion that Springfield is ready to rock. And it seems equally unacceptable to deny the suspicion on the grounds that we don't know the preparedness *vis-à-vis* rocking of our fellow concert goers:

(29) a. Rock Star: Is Springfield ready to rock?
b. Crowd: ??We don't know!

Confirming such a collective query thus seems to reveal two things. First, just as in the case of BEMs, there is a reasonable expectation that we can apply something very much like DEFEASIBLE CLOSURE: if we think that each of the group is ready to rock, we may infer defeasibly that the group is itself ready to rock. Second, that the group members can *anticipate* that the base-level facts for the other group members do in fact obtain, that all the others are ready to rock.

There are three especially noteworthy features we want to point out. First, this represents a bending of normal conversational rules—it's not the norm of assertion we are living up to here when we all shout *Yes!*. But if we couldn't be relied on to bend them in this way, the rock star's question would be an infelicitous (and not just hackneyed) way to begin the show. Second, just as with BEMs, we could all get it wrong. The entire audience could (prospectively) live up to our linguistic and epistemic duties and answer *Yes!* and have it turn out that, despite our enthusiasm, we were not *qua* group ready to rock. And third, if the two of us *qua* individuals are not ready to rock, then we can conclusively and without any anticipation about the readiness of our fellow audience members answer the rock star's query with a *No*. If we aren't ready, then no group to which we belong could be ready.

Being in a conversation in which BEMs are issued is a lot like going to a rock concert. The strong group readings in our example are floated or put in play by Alex. That amounts to something like a collective conjecture and we say that the same type of conversational rules can apply for taking up such readings as apply in answering

whether Springfield is ready to rock. That means that hearers can confirm a BEM and take up one of its strong group readings even if they are not in a position to flat out assert such a strong disambiguation. Denying a BEM does not involve the same bending of normal conversational rules. But neither does answering the rock star's collective query negatively: if we are not ready to rock, then it just is not true that Springfield is, and we may say so. So it is with BEMS: if the prejacent is not compatible with what a hearer knows then it is not compatible with what any group to which she belongs knows.

So we can pair DEFEASIBLE CLOSURE with another pragmatic principle, one targeted to the strong readings of BEMS that can be floated:

(30) ANTICIPATION
Suppose an utterance of $might(B)(\varphi)$ by S at i puts in play the propositions $P_1, P_2, \ldots$. And suppose that these quantify over the information available to $G_1, G_2, \ldots$ respectively. Then if $H \in G_i$ and $f_H(i) \cap [\![\varphi]\!] \neq \emptyset$, then H may infer defeasibly that $f_x(i) \cap [\![\varphi]\!] \neq \emptyset$, for each $x \in G_i$.

This is, of course, defeasible and need not be appropriately exploited in every context. All we require is that, like with DEFEASIBLE CLOSURE, it is suitably conventionalized and that conversational partners can expect each other to appeal to it when reasonably appropriate. But even if we do appeal to ANTICIPATION, we may have to retract the judgments based on it. We are inclined to answer *Yes!* when asked if Springfield is ready to rock. But if we hear booing coming from the audience, we will retract that. Similarly, if Investigator #2 pipes up with information to the contrary:

(31) a. Investigator #2: No, the window was locked from the inside.
b. Investigator #1: Oh, OK. He can't have got out that way.

This kind of retraction plays a big role in training CIA agents. No sense can be made of this, they say, unless truth values of sentences involving BEMS are sensitive to contexts and indices and points of assessment to boot. Note that when Investigator #1 retracts in this way, she retracts the strong reading floated. If the issue of the day had been what she knew and not how the culprit got away, she could stick to her guns just fine. That proposition is of little use now, but isn't and wasn't false. So there is no more reason to think this behavior of BEMS in conversation points in the direction of a radically relativized semantics than does finding out that, *contra* our expectations, Springfield is not ready to rock. It is perhaps reason for sadness, for rocking is what Springfield ought to want to do, but no reason for despair (or signing up with the CIA).

6. Eavesdroppers

The literature on readings of BEMS has established, as we mentioned earlier, that BEMS can be interpreted as making claims about evidence beyond the current awareness of speaker and hearer. Concomitantly, we will say that in particular contexts there may be many more than just a few readings for a BEM that are put in play.

The eavesdropper cases that provided much grist to the CIA mills are relevant here. Imagine that Chuck, unbeknownst to Alex and Billy, is monitoring them as they're trying to find the keys. When Alex says *They might be in the car*, Chuck says to himself (or to us who are watching this particular morality play) *She is wrong. They can't be in the car because I saw Billy come into the house with them*. What justifies Chuck's rejection of the BEM? Within the logic of our analysis, Chuck's rejection of the BEM can only be felicitous if he is part of a *G* that is within the relevant cloud of contexts. In other words, Chuck can felicitously reject the BEM as long as the BEM in Alex's mouth had as one of its multiple readings a reading where it was a claim about the information state of a group to which Chuck belongs.

Now, can that be? Can it be that in Alex's mouth the BEM had as one of its multiple readings a reading where it was a claim about the information state of a group to which Chuck belongs? It is by now routine to spin such scenarios in a way that makes Chuck someone who is completely unknown to Alex and Billy. So, how can Alex's BEM have a reading where its claim hinges (partly) on what Chuck knows? Simple: imagine that what makes someone part of a relevant group for a BEM is that they are engaged (in some sense) in the same investigation as the overt partners in the conversation. CIA agents complain that this makes BEMs too strong to allow any speaker to assert them. We agree. It is extremely unlikely that Alex is asserting the BEM under the reading that includes Chuck. But that is not what our story says. Alex is licensed to utter the BEM as soon as she is in a position to assert it under one relevant reading (the solipsistic one, typically). But since she utters it as a *bare* epistemic modal, she thereby puts into play multiple readings and it is one of those that Chuck rejects, which then results in the prejacent being rejected as well.

We believe there are limits to what constitutes a relevant group involved in the investigation. Here are two examples from our "CIA Leaks" that make the point. Suppose we are putting a randomly chosen card in an envelope. You catch a glimpse of the card and know that it is a black-suited face card. You say (32–a). Then, ten years later when we open the envelope—it's the Jack of Clubs—we cannot complain with (32–b):

(32) a. You: It might be the King of Spades.
b. Us [ten years later]: ??Wrong!/What you said is false!

Or consider the case of Detective Parker. He has been going over some old transcripts from Al Capone's courtcase in the 1920s—Capone is being asked about where some money is in relation to a particular safe:

(33) a. Capone: The loot might be in the safe.
b. Parker: ??Al was wrong/What Al said is false. The safe was cracked by Geraldo in the 80s and there was nothing inside.

There is just no relevant sense in which we-in-ten-years are involved in the same investigation that you-now are in. And there is no relevant sense in which the investigation that Capone and the DA were party to is now the same one taken up by Parker.

But of course there are borderline cases. If it's borderline whether a hearer *H* is part of a relevant group—borderline whether any group she belongs to is relevant—then we would expect the speaker's intuitions to be equivocal about attributions of falsehood or error. Since, in fact, we think there is such variability in intuitions, that makes us pretty happy.

7. Conclusion

We submit that our view of BEMs has clear advantages over other currently fashionable approaches. In particular, it involves no innovations in the semantics of epistemic modals. We make do with a standard context-dependent semantics. The new claim is that the strange properties of BEMs derive from the fact that their indeterminacy is used by design to achieve spectacular results.

References

Dalrymple, M., Kanazawa, M., Kim, Y., Mchombo, S., and Peters, S. (1998). "Reciprocal expressions and the concept of reciprocity", *Linguistics and Philosophy* 21(2): 159–210.

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

van Deemter, K. (1998). "Ambiguity and idiosyncratic interpretation", *Journal of Semantics* 15(1): 5–36.

Egan, A. (2007). "Epistemic modals, relativism, and assertion", *Philosophical Studies* 133(1): 1–22.

——Hawthorne, J., and Weatherson, B. (2005). "Epistemic modals in context", in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 131–70.

Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1995). *Reasoning about Knowledge* (Cambridge, Mass: MIT Press).

von Fintel, K., and Gillies, A. S. (2007). "An opinionated guide to epistemic modality", in T. S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology: Volume 2* (Oxford: Oxford University Press), 32–62.

—— —— (2008). "CIA leaks", *Philosophical Review* 117(1): 77–98.

—— —— (2010). "*Must* . . . stay . . . strong!", *Natural Language Semantics* 18(4): 351–83.

Hacking, I. (1967). "Possibility", *Philosophical Review* 76(2): 143–68.

Kratzer, A. (1977). "What *must* and *can* must and can mean", *Linguistics and Philosophy* 1(3): 337–55.

—— (1981). "The notional category of modality", in H.-J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*, number 6 in Research in Text Theory, (Berlin: de Gruyter): 38–74.

—— (1986). "Conditionals", *Chicago Linguistics Society*, 22(2): 1–15.

—— (1991). "Modality", in A. von Stechow and D. Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, (Berlin: de Gruyter), 639–50.

Lackey, J. (2007). "Norms of assertion", *Noûs* 41(4): 594–626.

MacFarlane, J. (this volume). "Epistemic Modals Are Assessment-Sensitive", ch. 5.

Poesio, M. (1996). "Semantics ambiguity and perceived ambiguity", in K. van Deemter and S. Peters (eds.), *Semantic Ambiguity and Underspecification* (Stanford, CA: CSLI) 159–201.

Raskin, V. (1985). *Semantic Mechanisms of Humor* (Dordrecht: Reidel).

Schwarzschild, R. (1999). "How specific is an utterance? Comments on Glanzberg (1999)." Handout from a talk given at the conference on "Mapping the Semantics-Pragmatics Boundary: Context-Dependence", Cornell University, March 26–8, 1999.

Stalnaker, R. (1978). "Assertion", in P. Cole (ed.), *Syntax and Semantics*, Vol. ix (New York: Academic Press), 315–32.

Stephenson, T. (2007a). "Judge dependence, epistemic modals, and predicates of personal taste", *Linguistics and Philosophy* 30(4): 487–525.

——(2007b). *Towards a Theory of Subjective Meaning*, Ph.D. thesis, Massachusetts Institute of Technology.

Su, S. P. (1994). *Lexical Ambiguity in Poetry* (Studies in Language and Linguistics; London: Longman).

Teller, P. (1972). "Epistemic possibility", *Philosophia* 2(4): 302–20.

Weiner, M. (2005). "Must we know what we say?", *Philosophical Review* 114(2): 227–51.

Williamson, T. (1996). "Knowing and asserting", *Philosophical Review* 105(4): 489–523.

Winter, Y. (2001). "Plural predication and the strongest meaning hypothesis", *Journal of Semantics* 18(4): 333–65.