DRAFT
<fintel@mit.edu>
<thony@umich.edu>

'Might' Made Right*

Kai von Fintel Anthony S. Gillies

Draft: July 2007

1 Introduction

Modals — *might*, *must*, *possibly*, *necessary*, *have to*, *can*, *ought to*, *presumably*, *likelier* and the rest — are quantifiers over a domain of possibilities. Different flavors of modality then 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*. More generically, they quantify over possibilities compatible with the information at hand. As we will sometimes put it: they quantify over what's not ruled out by a relevant *information state*.

But *whose* information state is it that determines the set of possibilities an epistemic modal claim quantifies over? A tempting answer — what we take to be, plus or minus a bit, the canon — is that the context decides. That would make these modals context-dependent in a pretty thorough-going way and also make them rather flexible in their range of possible readings.

The kind of context-dependence embodied by the canon has recently gone in for some rough treatment. CIA agents — those arguing that epistemic modals

* 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 [7, 8]. 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, 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.

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.[1] (Relativism is the new black.) It is not a state secret that we have our doubts about whether CIA theories hold water, and so (for the most part) we will be setting them to one side here.

Still, the canon is not quite right. But it is not all wrong either. What we want to do here is show how the canonical context-dependent semantics for epistemic modals can be supplemented with a picture of the pragmatics and conversational dynamics of these modals.

Just what information state does an epistemic modal in context quantifier over? We will sketch an answer along these lines. Bare epistemic modals — "bare" because they occur without accompanying restrictors — quantify over the information available to a contextually relevant group of investigators. But, given a context, there are multiple ways of drawing the group boundaries. 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. But that indeterminacy is — as we will say — indeterminacy by design.

2 Bare Epistemic Modals

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

(1) $\textsc{modal}(B)(\varphi)$

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.[2] 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

1 Some recent CIA agents: Egan [4], Egan et al. [5], MacFarlane [16], Stephenson [21, 22].

2 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 preceeding 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.

domain of possibilities determined by B, and we get different modals.[3]

One way (the value of) B gets determined is as the denotation of restricting phrases like *in view of*.[4] 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 the 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.

We call such (occurrences of) modals *bare modals*. 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: 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.

3 This is the now canonical analysis of the logical structure of modal statements developed by Kratzer [11, 12, 14].

4 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) $[\![B]\!]^{c,i}$.

3 Possible Readings

3.1 Flexibility

The canon has it that BEMs like *might* and *must* are quantifiers over information states, true at a context–index pair just in case the prejacent is true in some/every possibility compatible with the contextually-relevant information state. Given the flexibility of what gets to counts as relevant 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.

Still, even once we add some bells and whistles, the canon isn't quite right. The very flexibility that seems its greatest asset — the range of meanings it straightaway yields — requires contexts to do quite a lot of work that they just don't seem to do. Before getting to that, though, we want to tell our preferred version of the canon-cum-bells-and-whistles.

The range of possibly contextually relevant information states that determine the set of possibilities quantified over by BEMs 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 [13]:

> 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 [3] points out — it would never make any sense to say *I don't know whether might p*. But it can. DeRose's example (adapted just a bit): 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.[5]

So, 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.

Here, there are three issues. First: how do we aggregate the information states of individual agents that make up a group? Second: how far can group membership be stretched — can it involve eavesdroppers? And third: is non-solipsism enough objectivity? We are going to return to the second question a bit later, once our pragmatic story has been told in enough detail to see how we can deploy it for the case of eavesdroppers. The other two issues we'll take on here.

---

5 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.

3.2 Objectivity

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.

Or imagine a salvage crew searching for a ship that sank a long time ago.[6] The mate of the salvage ship works from an old log to determine where the dive team should be deployed. But he makes some mistakes in his calculations, and concludes that it is compatible with the information in the log that the wreck is in a certain bay. And so he announces to the crew:

(9) The hulk might be in these waters.

No one knows anything to the contrary. But in fact, as it turns out later, it simply was not possible for the vessel to be in that bay — more careful examination of the log would show that the boat must have gone down at least thirty miles further south.[7]

A natural reaction to this sort of case — the reaction that Hacking and DeRose push for — is to insist that it's not quite what the members of the group know that counts but what they can come to know, what is (borrowing Egan's [4] nice phrasing) within their "epistemic reach". Of course, then we have to decide how to constrain the notion of epistemic reach so that the story does not unravel. 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

6 The example is due to Hacking [10].

7 We do note, however, that intuitions about this kind of case aren't uniform.

of the matter about φ and yet *might* φ and *might not-*φ both seem true. An example from Teller [24]:

(10) 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 fares 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:

(11) 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.[8]

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 behind curtains or in the car. The project of gerrymandering epistemic reach to fit these boundaries has all the hallmarks of a project we don't want to take up. So we will 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 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

8 Another example (an elaboration of one in [9]): the forecast called for a 50% chance of rain. We are inside, and the curtains — even those just two inches off your elbow — are drawn. You are about to head out in your expensive cashmere (can't get it wet). A warning is issued:

(i) Take an umbrella — it might be raining out.

This seems a fully appropriate and true thing to say (even if it turns out that it isn't raining just now) even though there is a contextually relevant way of coming to know the fact of the matter about the current weather.

aren't in the car — cannot.[9]

3.3 Aggregation

That still leaves us with the question of *how* we aggregate the information states of a contextually determined group of agents. 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.

Now, then, some options. We might — as Hacking and DeRose both do — say that a BEM requires that no member of 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:

(12) $\quad [\![\mathit{might}\ \varphi]\!]^{c,i} = 1$ iff $\forall x \in G_c\colon\ \exists w \in f_x(i)$ 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.

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 $[\![B]\!]^{c,i}$:

(13) $\quad [\![\mathit{might}(B)(\varphi)]\!]^{c,i} = 1$ iff $\exists w \in [\![B]\!]^{c,i}:\ [\![\varphi]\!]^{c,w} = 1$

What is left 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

9 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 our paper "An Opinionated Guide to Epistemic Modality" [8].

the members of a group know.[10] A first pass is to take what everyone knows, simply collecting the possibilities compatible with what each agent knows:

$$\mathsf{E}_G(i) = \bigcup_{x \in G} f_x(i) \tag{14}$$

We could then supplement the skeletal semantics by identifying $[\![B]\!]^{c,i}$ with $\mathsf{E}_{G_c}(i)$. But this would be too weak a modal base to be plausible: a BEM *might* φ would be true in any context provided some relevant agent doesn't know that φ is false. That is too many *might*s coming out true and would predict that the dual *must*s are harder to be true than they in fact are.

Even further along that end of the spectrum (where it takes a lot for a group to know something and thus very little for a *might*-statement to be true) lies the common ground — what is common knowledge — between the group members. It is a bit easier to see what this amounts to by working not with selection functions, but accessibility relations. Where x is an agent with selection function f_x, let $R_x = \{\langle w, v\rangle \colon v \in f_x(w)\}$. Given an agent x and another y, we can take the (relational) composition of their relations: $R_x \circ R_y$. This will be the set of pairs $\langle w, v\rangle$ such that there is an intermediate possibility linking the two — a u such that u is compatible with what x knows at w and such that v is compatible with what y knows at u. The common ground among the group G is exactly the set of possibilities that can be linked in this way from i by composing the relations for the agents in it:

$$\mathsf{C}_G(i) = \{w \colon \exists x_1, \ldots x_n \in c \colon \langle i, w\rangle \in R_{x_1} \circ \cdots \circ R_{x_n}\} \tag{15}$$

Taking BEMs as quantifies over the common ground would be to combine (15) with the following proposal for aggregating a modal base, identifying $[\![B]\!]^{c,i}$ with $\mathsf{C}_{G_c}(i)$.[11] But since common knowledge is comparatively harder to come by, the set of worlds compatible with it is pretty large, even larger than in the case of $\mathsf{E}_{G_c}(i)$. And so existential quantifiers over it would be too easily true to count this as a plausible thing that *might* might mean.[12]

---

10 The classic reference for multi-agent epistemic logic is Fagin et al. [6].

11 Equivalently: the denotation of B, with respect to c and at i, is the set of possibilities related by i in the reflexive transitive closure of the R_x's. The detour through accessibility relations — both in this form and in the version in (15) — is, of course, inessential. But it does simplify by allowing us to just consider the closure (respectively: the composition) of the relations in a straightforward way.

12 *Might* ϕ would be true, for example, in a case where everyone knows that not ϕ but someone (x) in the group doesn't know that everyone in the group knows that x knows that not ϕ.

More plausible given the actual behavior of BEMs is to *pool* what the individual agents know. We do this by intersecting the possibilities compatible with each information state at i:

(16) $\mathsf{D}_G(i) = \bigcap_{x \in G} f_x(i)$

It is possible for the group to have the information that rules out the prejacent without any of the group members having that information. Indeed, that is why the notion is so useful in logics modeling the behavior of distributed systems.[13] And the supplement to the skeletal semantics is simply to identify the possibilities determined by the modal base with respect to c and i with what is distributed among the relevant agents at i: $[\![B]\!]^{c,i} = \mathsf{D}_{G_c}(i)$.[14]

A feature about distributed information that plays a role for us is simply this: whenever one group is included in another, whatever is distributed information in the first must also be distributed information in the second. That is, the following always holds:

(17) If $G \subseteq G'$ then, for any i, $\mathsf{D}_{G'}(i) \subseteq \mathsf{D}_G(i)$

Thus whatever is true throughout what the smaller group knows, is true throughout what the bigger group knows.

It is tempting to think that BEMs quantify over the pooled information in a group and that the context determines what the relevant group is. That would give an easy explanation of the cases above, and by being egalitarian about agent-hood we can say why stores of information in ships logs and computers count to whether a BEM is true and why rain behind curtains and keys (not in) a car don't. And it makes it easy to see how there could be such a distribution of readings available: by varying the group, we vary what might and might not be given that group.

---

13 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.

14 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 (16) for non-modals, 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 (16), 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).

This way of putting the canon — that is to say our dressed up version of the canon — has some noteworthy features. We'll mention three. First, if $G = \{\text{speaker}\}$ then $\mathsf{D}_G(i) = f_s(i)$. Thus with respect to a context that determines such a G as the c-relevant group, the modals end up being solipsistic expressions of what is compatible with what the speaker knows after all. Second, there is nothing in any of this that requires a hardwired commitment that the speaker is always a member 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 we have 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. Third, we noted in passing that in some cases — notably in our Schmolmes example and similar Hacking-type salvage ship cases — there is some unexpected divergence of judgments. Some judge the relevant BEMs as straightforwardly false; others disagree. And some just get queazy. This distribution of judgments falls out rather nicely as a consequence of the relative level of comfort speakers have in treating inanimate stores of information as agents — the more egalitarian among us find the BEMs in these cases flatly false and chauvinist-dialect speakers swallow them more easily.

Since it is tempting to think that BEMs quantify over the pooled information in a group and that the context determines what the relevant group is, and since we like to give into 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.

4 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:

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

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:

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

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

(20) Billy: No, I still had them when we came into the house.[15]

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. Call this speaker-centric resolution $c_s = \{\text{Alex}\}$, and 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. Call this groupwise resolution $c_g = \{\text{Alex}, \text{Billy}\}$, and the associated reading the $A{+}B$-reading. Here the BEM is about the information that the non-trivial group of Alex and Billy have. 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.[16]

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 (20). 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

15 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 how to exactly derive this obligation, even if it seems blindingly obvious.

16 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.

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.[17] But the reply in (19), *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 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

(21) 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 (21) 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 at i, Billy doesn't know where the keys are, then her reply in (19) would be on the mark since in this case there would be keys-in-the-car possibilities in $\mathsf{D}_{c_g}(i)$.[18] And if, at i, Billy knows they're not in the car, then her private knowledge — the fact that there are no keys-in-the-car worlds in $f_B(i)$ — is sufficient to guarantee that there are no such keys-in-the-car possibilities in $\mathsf{D}_{c_g}(i)$. The denial in (20) 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 (18) in the first place. She does not seem to be within her linguistic rights to be claiming that the group cannot rule out the prejacent.[19] After all, Alex does not know

17 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.

18 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.

19 As far as we know, the earlier literature on group readings of BEMs [3, 24] did not discuss the question of what justifies the assertion of such a reading by a speaker who does not have

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[20] 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 of the predicted readings is the one that the context determines as the interpretation of the BEM that Alex uttered.[21]

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 in "CIA Leaks" [7] why not. 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.

---

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*.

20 We don't pick sides in the fight, see Williamson [26], Weiner [25], and Lackey [15] for some of the contestants.

21 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.

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.

5 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.[22]

This, we admit, sounds a little exotic. But, in fact, it is no more exotic than the common chimera. Not all actual contexts of use for a natural language sentence will successfully resolve all indeterminacies and ambiguities of that sentence. 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:

(22) 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:

(23) 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).[23] And yet other times, the

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

23 The example in (23) is due to van Deemter [2]. Idealizing, we might say that Bond is

indeterminacy is there as if by design, and this is what we claim is going on with BEMs.[24]

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. The context of the conversation really does not provide a determinate resolution and 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. We begin by saying what *travels* in a conversation, or what proposition(s) the speaker of a BEM *puts into play*:

---

employing the strategy of diagonalization [20] here: he computes the proposition that whoever the guide is referring to is trying to shoot him.

24 It has been observed that sometimes ambiguous sentences are used purposefully to convey more than one proposition, namely in the context of jokes [18], poetry [23], and other less than straightforward uses of language; see Poesio [17] for discussion. Our proposal here is that a mutiplicity of meaning is also detectable in at least ostensibly more straightforward uses of language.

(24) 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 $\mathit{might}(B)(\varphi)$ with respect to C at t puts into play the set of propositions P such that for some $c \in C$: $[\![\mathit{might}(B)(\varphi)]\!]^c = P$.

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.[25] 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} and {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 (linguistic and epistemic) 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

25 Poesio [17], for example, gives a semantics for an underspecified language that assigns sets of standard meanings to expressions of that language.

speaker utters an underspecified sentence she must be — deontic *must*, that is — 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:

(25) ASSERT:
Suppose an utterance of $might(B)(\varphi)$ 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. Obviously, the minimal requirement if the group containing just the speaker is among the the candidate groups then is that the speaker needs to be justified in asserting the A-reading. That will be the weakest reading she will stand in the appropriate relation to. In other words, the speaker 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.

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 say that the hearer should be guided by what response to which proposition will be most informative in the conversation. This will in fact lead to a dominance of negative replies.

(26) 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 the right thing to do 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, we think that 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.[26]

26 Instead of justifying our principle in (26) 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 [1, 27].

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 distributed 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 actually is 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. 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. We propose that there is a defeasible principle as follows:

(27) 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, etc. 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.

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 the speaker 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: the 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.[27]

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:

(28) 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 there *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:

(29) Every student was at the meeting.

---

27 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.

It does not seem implausible to say that the precise delimitation of the contextual domain of quantification for (29) can often be indeterminate in a realistic context. We suggest that the speaker of (29) in such a context has put into play a set of propositions 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:

(30) 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 perhaps almost always irrelevant, except in cases of antagonistic conversations.

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.

6 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 (A+B-, A+C-, and A+B+C-readings).[28]

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 (26) 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 (28), 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

28 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*.

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:

(31) 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?

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:

(32) 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:

(33) 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 quite 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:

(34) 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 distributed among $G_1, G_2, \ldots$ 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:

(35) 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 context and indices and points of assessment to boot. But 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).

7 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 (36a). Then, ten years later when we open the envelope — it's the Jack of Clubs — we cannot complain with (36b):

(36) 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 1920's — Capone is being asked about where some money is in relation to a particular safe:

(37) 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 80's 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 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.

8 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.

Bibliography

[1] Dalrymple, Mary, Makoto Kanazawa, Yookyung Kim, Sam Mchombo, and Stanley Peters (1998). Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy* 21(2):159–210, doi:10.1023/A:1005330227480.

[2] van Deemter, Kees (1998). Ambiguity and idiosyncratic interpretation. *Journal of Semantics* 15(1):5–36, doi:10.1093/jos/15.1.5.

[3] DeRose, Keith (1991). Epistemic possibilities. *The Philosophical Review* 100(4):581–605, doi:10.2307/2185175.

[4] Egan, Andy (2007). Epistemic modals, relativism, and assertion. *Philosophical Studies* 133(1):1–22, doi:10.1007/s11098-006-9003-x.

[5] Egan, Andy, John Hawthorne, and Brian Weatherson (2005). Epistemic modals in context. In Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, 131–170, Oxford: Oxford University Press.

[6] Fagin, Ronald, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi (1995). *Reasoning about Knowledge*. MIT Press.

[7] von Fintel, Kai and Anthony S. Gillies (2007). CIA leaks, URL http://mit.edu/fintel/cia_leaks.pdf, final ms, MIT and University of Michigan, to appear in *The Philosophical Review*.

[8] von Fintel, Kai and Anthony S. Gillies (2007). An opinionated guide to epistemic modality, URL http://mit.edu/fintel/fintel-gillies-ose2.pdf, ms, MIT and University of Michigan, to appear in Tamar Gendler Szabó and John Hawthorne (eds.) *Oxford Studies in Epistomology, Vol.2*, Oxford University Press.

[9] Gillies, Anthony S. (2001). A new solution to Moore's paradox. *Philosophical Studies* 105(3):237–250, doi:10.1023/A:1010361708803.

[10] Hacking, Ian (1967). Possibility. *The Philosophical Review* 76(2):143–168, doi:10.2307/2183640.

[11] Kratzer, Angelika (1977). What *Must* and *Can* must and can mean. *Linguistics and Philosophy* 1(3):337–355, doi:10.1007/BF00353453.

[12] Kratzer, Angelika (1981). The notional category of modality. In H. J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts. New Approaches in Word Semantics*, 38–74, Berlin: de Gruyter.

[13] Kratzer, Angelika (1986). Conditionals. *Chicago Linguistics Society* 22(2):1–15.

[14] Kratzer, Angelika (1991). Modality. In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, 639–650, Berlin: de Gruyter.

[15] Lackey, Jennifer (2007). Norms of assertion, URL http://www.niu.edu/phil/~lackey/Norms%20of%20Assertion-Nous-Revised.pdf, ms, to appear in *Noûs*.

[16] MacFarlane, John (2006). Epistemic modals are assessment-sensitive, URL http://sophos.berkeley.edu/macfarlane/epistmod.pdf, ms, University of California, Berkeley, forthcoming in an OUP volume on epistemic modals, edited by Brian Weatherson and Andy Egan.

[17] Poesio, Massimo (1995). Semantics ambiguity and perceived ambiguity. In Kees van Deemter and Stanley Peters (eds.), *Semantic Ambiguity and Underspecification*, Stanford, CA: CSLI.

[18] Raskin, Victor (1985). *Semantic Mechanisms of Humor*. Dordrecht: Reidel.

[19] Schwarzschild, Roger (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 26-28, 1999.

[20] Stalnaker, Robert (1978). Assertion. In Peter Cole (ed.), *Syntax and Semantics*, volume 9, 315–322, New York: Academic Press.

[21] Stephenson, Tamina (2007). Judge dependence, epistemic modals, and predicates of personal taste, URL http://semanticsarchive.net/Archive/mU2OGM3Z/, ms, MIT, to appear in *Linguistics and Philosophy*.

[22] Stephenson, Tamina (2007). *Towards a Theory of Subjective Meaning*. Ph.D. thesis, Massachusetts Institute of Technology.

[23] Su, Soon Peng (1994). *Lexical Ambiguity in Poetry*. Studies in Language and Linguistics, London: Longman.

[24] Teller, Paul (1972). Epistemic possibility. *Philosophia* 2(4):302–320, doi:10.1007/BF02381591.

[25] Weiner, Matthew (2005). Must we know what we say? *The Philosophical Review* 114(2):227–251, doi:10.1215/00318108-114-2-227.

[26] Williamson, Timothy (1996). Knowing and asserting. *The Philosophical Review* 105(4):489–523, doi:10.2307/2998423.

[27] Winter, Yoad (2001). Plural predication and the strongest meaning hypothesis. *Journal of Semantics* 18(4):333–365, doi:10.1093/jos/18.4.333.