

CIA Leaks

Kai von Fintel

MIT

Anthony S. Gillies

Michigan

1. The Plot

Some seemingly ordinary observations about the epistemic modals *might* and *must* have recently led some to an extraordinary measure. The extraordinary measure is to insist that sentences like

- (1) Billy might be at the party

only get assigned truth-values relative to contexts of utterance, indices of evaluation, and—the new wrinkle—points of assessment. We dub such semantic analyses CIA theories.¹ Our goal here is to catalogue some central problems that any CIA theory must solve. We begin by briefly sketching a motivation for CIA theories and a generic formulation of them. Then we lay down our gauntlet, showing where and why there are troubling CIA leaks. We end by outlining an alternative semantics that we develop elsewhere.

For comments, we would like to thank Josh Dever, John Hawthorne, Angelika Kratzer, two anonymous *Philosophical Review* referees, 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, and at the 2006 Pacific APA.

1. A not-wholly misleading handle is to think of points of assessment as the *judge* or *point of view* relative to which truth-values are assigned. Some recent CIA agents: Egan, Hawthorne, and Weatherson (2005); Lasersohn (2005); MacFarlane (2004); Stephenson (forthcoming).

Philosophical Review, Vol. 117, No. 1, 2008

DOI 10.1215/00318108-2007-025

© 2007 by Cornell University

2. CIA Theories

The canon about *might* and *must* is that they are context-dependent quantifiers over possibilities. But the quantifiers are not unrestricted, and so the modals are not absolute. So (1) has a logical form along the lines of

$$(2) \quad \textit{might}(B)(\varphi)$$

where B restricts the domain over which *might* quantifies. Epistemic modals quantify over possibilities compatible with *what is known*—more generally, those possibilities compatible with the *evidence available* or the *information at hand* in a context.

Formally the restricting gets done by assigning (in a context c , at an index i) as denotation to B a set of worlds:²

$$(3) \quad \llbracket B \rrbracket^{c,i} = \{v : v \text{ is compatible with the } c\text{-relevant information at } i\}$$

A modal in a context c at i then quantifies over $\llbracket B \rrbracket^{c,i}$:

$$(4) \text{ a. } \llbracket \textit{might}(B)(\varphi) \rrbracket^{c,i} = 1 \text{ if and only if } \exists w' \in \llbracket B \rrbracket^{c,i} : \llbracket \varphi \rrbracket^{c,\langle w', t_i \rangle} = 1$$

$$\text{ b. } \llbracket \textit{must}(B)(\varphi) \rrbracket^{c,i} = 1 \text{ if and only if } \forall w' \in \llbracket B \rrbracket^{c,i} : \llbracket \varphi \rrbracket^{c,\langle w', t_i \rangle} = 1$$

(In what follows, we will sometimes refer to φ as the *prejacent*, a useful term introduced by our medieval colleagues.) If it is only the speaker's information that counts in c , then $\llbracket B \rrbracket^{c,i}$ will be the set of worlds compatible with the evidence she has access to at i . Analyses in the spirit of the canon predict context sensitivity of *might* and *must*: since what is known can vary with context—either by varying who the relevant knowers are or by varying what it is they know—so can the domains over which the modals quantify.

Given both the power and the flexibility of the canon, it is a bit of a surprise that simple cases of modal disagreement seem to pose a problem for it. Here is an old example from Kratzer (1986):

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.

2. B is often called the *modal base*, following Kratzer's work on modality (see, for example, Kratzer 1981). We will take indices to be world-time pairs.

In view of *your* evidence, it cannot possibly be Fred, it must be Martin. If this is so, *my* utterance of (5) and *your* utterance of (6) are both true.

- (5) The person approaching might be Fred.
- (6) The person approaching cannot be Fred.

Had *I* uttered (6) and *you* (5), both our utterances would have been false.

What Kratzer does not comment on is the feeling that “you” and “I” are disagreeing: (5) and (6) seem incompatible. Kratzer’s verdict is that (5) and (6) are true *in their contexts*. If (5) in her mouth expresses the proposition that the person approaching is Fred is compatible with what *she* knows and (6) in your mouth expresses the proposition that the person approaching is Fred is not compatible with what *you* know, then you have no beef with Kratzer. But if you *are* disagreeing—you affirm a proposition she denies—then you cannot both be saying something true.

MacFarlane (2004) discusses a somewhat similar case:

- (7) a. Sally: Joe might be in Boston.
- b. George: He can’t be in Boston. I saw him in the hall five minutes ago.
- c. (i) Sally: Oh, then I guess I was wrong.
- (ii) Sally: Oh, OK. So he can’t be in Boston. Nonetheless, when I said “Joe might be in Boston,” what I said was true, and I stand by that claim.

Although (7ci) seems a natural enough reply, (7cii) is pretty stilted. MacFarlane argues that this is problematic since if Sally’s *might*-claim in (7a) quantified over what she knew when she uttered it, then we would expect the reverse.³

And disagreement is not bound to the conversation: eavesdroppers can be pretty opinionated. Here is an example Egan (2007) makes some hay about. Bond and Leiter are in London, listening to a bug Bond planted in a conference room in SPECTRE’s headquarters in the Swiss Alps. Bond left behind some misleading evidence pointing to his presence in Zürich. Blofeld finds the evidence, takes it to be genuine, and turns to his second in command, Number 2:

- (8) Bond might be in Zürich.

3. The data here, as we will argue below, are neither innocent nor robust. But right now we are just trying to motivate the CIA theories a bit.

And Number 2 may well reply with a “That’s true.” But Leiter, hearing all this from London, is not at all inclined to say “That’s true” when *he* hears (8) from Blofeld, even though Leiter knows full well that it is compatible with everything that Blofeld knows—and indeed compatible with everything Blofeld’s conversational partners know—that Bond might be in Zürich. Instead, Leiter might say “Excellent: Blofeld’s wrong again!” or “That’s false.”

Cases of modal disagreement like these seem to pose a dilemma for the canon. On the one hand, each asserter of a *might*-claim is perfectly within their (linguistic and epistemic) rights to say what they said, and the thing they end up saying seems true. Given the facts about her utterance situation, (the early) Sally seems to be saying that Joe’s being in Boston is compatible with the facts she had at the time of utterance. And that seems true. Given the facts about his utterance situation, Blofeld seems to be saying that Bond’s being in Zürich is compatible with the facts he had at the time of utterance. And that seems true. But, on the other hand, we get the sense that there is a disagreement between you and Kratzer, between Sally and George, between Sally and her earlier self, and between Blofeld and Leiter. In each case, one of the parties assents to a *might*-claim that the other denies. It is hard to make room for disagreement within the simple family of semantic analyses. We could make the modal base more complicated—letting it select worlds compatible with the information of a wider pool of agents—but that, at first blush anyway, makes it harder to see how our asserters seemed perfectly within their rights to say what they said.

The CIA theory says we can dodge the choice: the semantics of epistemic modals is relative not only to a context and index, but also to an assessor—someone, somewhere, who is doing the interpreting. Information at the point of assessment—roughly, what the assessor knows—enters into the semantics, determining the possibilities that the modals quantify over. Let’s take a point of assessment a to determine an assessor or judge j_a , a world of assessment w_a , and a time of assessment t_a . Here is the most generic way of formulating CIA theories:

- (9) a. $\llbracket \textit{might}(B)(\varphi) \rrbracket^{c,i,a} = 1$ if and only if
 $\exists w' \in \llbracket B \rrbracket^{c,i,a} : \llbracket \varphi \rrbracket^{c,\langle w',t_i \rangle,a} = 1$
 b. $\llbracket \textit{must}(B)(\varphi) \rrbracket^{c,i,a} = 1$ if and only if
 $\forall w' \in \llbracket B \rrbracket^{c,i,a} : \llbracket \varphi \rrbracket^{c,\langle w',t_i \rangle,a} = 1$
 c. $\llbracket B \rrbracket^{c,i,a} = \{v : v \text{ is compatible with what } j_a \text{ knows at } t_a$
 in $w_a\}$

Usually w_a is the same as the world associated with the index (since we typically are in conversations only with worldmates), but t_a may well be different from t_i . When there is no risk of confusion, we identify points of assessment with their assessors.

The rough diagnosis that CIA theories have on offer is that Sally's initial claim is assigned true when she is the point of assessment and false when it is George (and her later self). Blofeld's utterance of (8) is true when he (and Number 2) is assessor and false when Bond and Leiter are. Insofar as we want to treat cases of modal disagreement as situations in which both parties speak truly but have substantive disagreement, CIA theories seem to offer a considerable overlap between our cake having and our cake eating.

We intend to rain on this parade, cake and all. And no one likes soggy cake.

3. Data on "Data"

A closer look at denial-correction sequences in discourse (and at their agreement-acceptance cousins) makes it clear that the motivating data for CIA theories is neither innocent nor unambiguous.

The first fact we want to draw attention to is that—*pace* the CIA—not all *mights* are retracted or rejected in the face of new evidence. Speakers can quite often *resist* the invitation to retract even if they have become better informed. Billy is looking for her keys. Alex is trying to help.

- (10) a. Alex: The keys might be in the drawer.
- b. Billy: (*Looks in the drawer, agitated.*) They're not. Why did you say that?
- c. Alex: Look, I didn't say they *were* in the drawer. I said they *might be* there—and they might have been. Sheesh.

And some constructions—for example, *mights* under *if*—are particularly resistant to rejection. We are wondering where John is, and you say

- (11) If John is not in his office, he might be in New York.

We go to his office and find a note saying "I'm in Boston." It is not natural for you to retract in this way.

- (12) ??I guess I was wrong.

Knowing what you *now* know, you could not truthfully utter (11), but that is no reason to invoke CIA theories.

The basic observation is that *solipsistic* readings for the modals—readings on which the modals quantify over the evidence available to the speaker at the time of utterance—are virtually always available. Sometimes, as in the case of Alex and Billy, those readings are made clear by conversational resistance: Alex sticks to her guns. And sometimes, as in the other kind of case, it is a certain construction that induces preference for the solipsistic resolution of the modal.

That brings us to the second observation we want to make. Even in the original disagreement and eavesdropping cases we glossed earlier, speaker intuitions are not anything like uniform. This, we say, is a fact that itself needs to be explained by our best theory. And CIA theories do not seem well equipped for this—assessorhood is pretty unambiguous, and that leaves little room for accommodating anything but the most hard-line retraction intuitions in those cases.

And now to the final particular worry we want to raise about the motivating data: denial (and acceptance) in discourse is pretty tricky business. The CIA theory gets its motivation from exchanges like

- (13) a. A: It might be that *P*.
 b. B: No, it can't be that *P*/No, it must be that not-*P*.

What is supposed to cause all the trouble for the canon is that if what *might* in *A*'s mouth means is a function of (only) her evidence, then although *B* denies what *A* says (*A* utters a sentence that *B* negates), they do not disagree—the proposition *A* expresses is not incompatible with the proposition *B* expresses.

But this is not decisive. To cut the kind of ice the CIA agents want cut, we need some reason to think that denials in discourse pretty much always track disagreement of the relevant kind. The facts, however, are more delicate than this. Consider:

- (14) a. A: I think it's raining out.
 b. (i) B: No, it isn't/No, it can't be.
 (ii) B: ??No, you don't.

B's felicitous denial (14bi) represents no disagreement with the proposition expressed by *A*'s utterance; (14bii) *does* represent that

disagreement, but is not felicitous. Denials don't always attach to the attitude claim (and sometimes cannot); the same goes for agreement.⁴

In the case of our modals, matters are worse: replies can target either the modal claim itself, or the prejacent—and with (apparently) equal effect. Pascal and Mordecai are playing Mastermind. After some rounds where Mordecai gives Pascal hints about the solution, Pascal says

(15) There might be two reds.

Mordecai, knowing the solution, has a range of possible responses:

- (16) a. That's right. There might be.
b. That's right. There are.
c. That's wrong. There can't be.
d. That's wrong. There aren't.

Before resorting to the CIA's extraordinary measures, we would like to see it carefully argued that flexibility in the target of denials and acceptances does not explain what needs explaining. A nice bonus would be some demonstration that the CIA's extraordinary measures leave room for this flexibility.

There is one more degree of possible divergence between the proposition expressed by an utterance and the actual target of a denial. A denial might not target the truth of the proposition expressed but the grounds for asserting that proposition. A *might*-claim is (pragmatically) more than just a profession of ignorance. By choosing a particular prejacent, the speaker is highlighting that possibility as one that should not be ignored.⁵ Further, there is often a reliance on positive evidence that makes that possibility seem to be a serious possibility. This means that a *might*-claim might be open to reproach, retraction, or disavowal, as soon as it becomes clear that the speaker was relying on assumptions that were not reliable—or, more generically, when it becomes clear that her epistemic state was not as it should have been—even though it is perfectly true that the prejacent was compatible with what the speaker

4. Two further complications: (i) denials can *target* presupposed, as opposed to asserted, material; (ii) presupposed material can also be the *source* of denial—when *A* says *P*, *B* can deny this (sometimes) by saying something that presupposes not-*P*.

5. Swanson (2006) lays out such a speech-act role for *might*-claims and builds a semantic theory around it. We are more inclined to stick to a standard semantics and to acknowledge that the pragmatic point of a *might*-claim is to highlight a possibility that should not be ignored.

knew.⁶ Blofeld could be said to be wrong not because his *might*-claim is false but because he is relying on misleading evidence, never a good move in his business.

The overarching worry we have about the CIA data is a worry about fit. On the one hand, it is relatively easy to construct cases in which we have pretty robust intuitions that: (i) a prior *might*-claim is perfectly acceptable; (ii) some intervening discourse reveals some new facts; and (iii) that same *might*-claim could not be acceptably uttered in the posterior situation. But all parties are—or should be—agreed that no pro-CIA conclusion follows from this. What is needed is evidence that the proposition expressed by the prior *might*-claim goes from true to false: and that is data that, roughly put, bridges the gap between denying an earlier claim and disagreeing with it.

To fill this gap, the intuitive data has to be pretty carefully cultivated. It is not enough to note that (sometimes) in posterior situations speakers will reject or retract prior *might*-claims. It is a particular sort of posterior retraction that is needed: the retractor has to make sure she is targeting the proposition expressed by the prior claim. That is why CIA agents tend to rely in the examples on retraction-talk like

- (17) a. What I said before is false.
 b. That is not true.

(Note that it is not *What I said was false* and the like that are relevant.) But about these kinds of cases—the ones that do the relevant work—speaker intuitions are far less robust.

Having expressed some worries about the motivating data, we now turn to our gauntlet.

4. Time Lag

According to the standard theory, *might* φ ends up quantifying over the worlds compatible with the relevant set of facts at world w_i at time t_i . That is, the modal base is determined relative to the index of evaluation i . In CIA-theories of epistemic modals, the modal base is determined relative to the context of assessment a instead.⁷ Now, let us look more closely at that claim. It is, formally speaking, open to CIA agents to take

6. Stephenson (forthcoming) raises similar considerations.

7. We note without further discussion that this makes epistemic modals special in a way that a more unified semantics for modals might not be happy with.

$\llbracket B \rrbracket^{c,i,a}$ to determine just about any set of worlds as a function of the point of assessment. We will focus on the most principled version that we can think of (not least because it makes for a theory that delivers predictions):

$$(18) \quad \llbracket B \rrbracket^{c,i,a} = \{v : v \text{ is compatible with } j_a \text{'s information} \\ \text{in } w_a \text{ at } t_a\}$$

This is more or less what we take to be the common core of the various CIA theories and what sets them apart from the canon.⁸

Points of assessment, equipped as they are with times, can be ordered according to the times associated with them. And, naturally enough, the *later* a point of assessment—that is, the later the time associated with a point of assessment—in general the *smaller* the set of possibilities compatible with it. Knowledge does tend to accumulate, after all. And as the set over which an existential modal quantifies gets *smaller*, it

8. A wrinkle: CIA theories can simply identify points of assessment with their assessors, making no room for a time parameter associated with that point that can differ from the time provided by the index (Lasersohn 2005; Stephenson, forthcoming). We thus get a modal base, at a c,i,a -triple, like this:

$$(i) \quad \llbracket B \rrbracket^{c,i,a} = \{v : v \text{ is compatible with } j_a \text{'s information at } i\}$$

That makes for a less radical CIA theory. But there is a cost to the conservatism. Consider a case of self-disagreement:

- (ii) a. Alex at t_1 : The keys might be in the car.
(... *Looks...* *They're not there.*)
- b. Alex at t_2 : Oh, I guess I was wrong. They can't be in the car.

If we identify points of assessment with judges, then it is what Alex knows at the time provided by the index that determines the modal base. But at t_1 she didn't (let's say) know that the keys weren't in the car. And so it looks like we get the distinctly non-CIA prediction that what Alex says at t_1 is true whenever Alex—even her later self—is the judge.

So there are two options for this would-be conservative strain within the CIA. First: say that judges are time-slices of individuals. That way the judge (or time-slice thereof) brings the time of assessment right along with her and overwrites the time associated with the index in determining the modal base. This makes for a full-blown, and no less radical, CIA theory—and the problems about time lags and tense shifting we lay out in this section and the next apply. Second: say that judges are individuals, full stop. This does make for a more conservative CIA theory and the time lag and tense-shifting problems do not apply to it. But now what needs explaining is the disagreement between Alex and her prior self, and that explanation—whatever it is—can't be what explains disagreement between parties and isn't what is operative in retraction. So, while conservative, the result is a much more fragmented theory than a full-blown CIA theory.

gets *harder* for such modals to be true. The CIA delivers a prediction: when $t_c < t_a$ it is in general the case that $\llbracket B \rrbracket^{c,i,a} \subset \llbracket B \rrbracket^{c,i,c}$. Thus, it is harder for simple *might*-claims to be true relative to an assessor in a later context of assessment a than it is for them to be true relative to an assessor in c , the original context of utterance. And the bigger the gap between t_c and t_a , in general the bigger the difference between the modal bases, and hence the easier it should be at a to reject *might*-claims issued at c . And the ease in rejection should increase as the gap between t_c and t_a grows.

But the facts go the other way: as t_a gets much later than t_c , it becomes increasingly silly to go in for the sort of rejection that the CIA predicts. 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 (19a). Then, ten years later when we open the envelope—it's the Jack of Clubs—we cannot complain with (19b)

- (19) 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 court case in the 1920s—Capone is being asked about where some money is in relation to a particular safe:

- (20) 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.

This time-lagging phenomenon would be unexpected were CIA theories on the right track. Instead, we should (ten years hence) think that what you said in uttering (19a) is false, and Parker should have no trouble in (20) disagreeing with (the long-dead) Capone.⁹

9. A CIA agent may well retreat to a weaker theory, claiming that not every point of assessment is a *legitimate* point of assessment—only *eligible* or *relevant* points count. Invoking relevant assessors, the CIA can then explain why we cannot disagree over large gulfs of time. To which we reply: no matter the notion of relevance, it will be determined by features of context not by the point of assessment. And we are the first to agree that limiting the extent to which CIA theories rely on the posited novel parameter reduces the implausible predictions such theories make.

5. Tense

The canonical context-dependence story about epistemic modals insists that the set the modals quantify over is determined relative to the index. So the semantic value of *might* φ at an index i , as we noted above, ends up depending on what happens at worlds compatible with the relevant set of facts at world w_i at time t_i . This makes for a pretty nifty division of labor when we consider shifty operators—like tense operators—that outscope the modals since it means that scope matters matter.

CIA theories instead say that the index is pretty much invisible to epistemic modals. This creates trouble once we consider sentences where the modal is embedded under a tense operator since whatever shifting of the time-coordinate that the tense operator does is promptly ignored by the modal.

Sophie is looking for some ice cream and checks the freezer. There is none in there. Asked why she opened the freezer, she replies:

- (21) a. There might have been ice cream in the freezer.
 b. $\text{PAST}(\textit{might}(\textit{ice cream in freezer}))$

It is possible for Sophie to have said something true, even though at the time of utterance she knows (and so do we) that there is no ice cream in the freezer.

But CIA theories do not deliver this prediction. Let P abbreviate *there is ice cream in the freezer*, and let t_c be the time of Sophie's utterance. Then $\llbracket \text{PAST}(\textit{might}(B)(P)) \rrbracket^{c, \langle w_c, t_c \rangle, a} = 1$ if and only if there is a recent-ish $t' < t_c$ such that $\llbracket \textit{might}(B)(P) \rrbracket^{c, \langle w_c, t' \rangle, a} = 1$. While the modal will pass on the shifted time to the evaluation of the prejacent, the modal base selects a set of worlds independent of the index. That is, the modal base of worlds we are quantifying over will be determined by what is known at a , in particular at t_a . So, since Sophie knows at t_a that there is no ice cream in the freezer, no worlds in $\llbracket B \rrbracket^{c, i, a}$ will be P -worlds. That means the embedded *might*-claim is bound to be false no matter the choice of t : $\llbracket \textit{might}(B)(P) \rrbracket^{c, \langle w_c, t \rangle, a} = 0$. And that means that the PAST-claim (21b) is predicted to be false.

The basic point is that CIA theories insist that the index of evaluation is pretty much invisible to, and irrelevant for, the interpretation of epistemic modals. This is conceptually pretty awkward: the CIA requires us to jettison the nifty way of dealing with embeddings in a general way. And it is also empirically pretty embarrassing: many a *might*-claim

outscooped by a past tense operator is predicted to be false (so long as at *a* the prejacent of the *might* is known to be false).

The other issues we want to raise are independent of issues involving the time-coordinates of indices and points of assessment, so we will generally suppress the *t*'s from now on.

6. Gibbarding

The next two kinds of problems involve various asymmetries between the evidence available to the speaker and the evidence available to the assessor. We first consider cases in which the assessor knows less than the speaker.

We will assume here that the CIA-analysis should extend to the dual of *might*: *must*. (If *might* and *must* are duals, it is hard to see how there is any room to favor CIA theories for one and not the other.) Here as well, the true CIA agent will adopt a semantics according to which what matters is what the assessor knows at the time of assessment. By the CIA's lights, an assessor j_a should reject any *must*-claim whose truth (with respect to *a*) would rule out worlds compatible with the facts at *a* (or whose truth would require such worlds to have already been ruled out).

CIA theorists tend to spill a lot of ink over cases of modal disagreement. That is not so surprising. What is surprising is how little attention is paid to cases of "uptake"—more or less plain vanilla information exchange between speaker and hearer.

Consider an analogue of the Gibbard phenomenon (Gibbard 1981).¹⁰ The Boss has two informants, Jack and Zack. There is a meeting of spies in a room, and The Boss, Jack, and Zack know that one and only one of their (conveniently named) comrades *P*, *Q*, *R* is a turncoat. Jack looks through his peep hole and sees clearly that it is either *P* or *Q* who is the turncoat, and Zack looks through his peep hole and sees clearly that it is either *Q* or *R* who is the turncoat. Each slips The Boss a note informing him:

- (22) a. [From Jack]: It must be that either *P* is the turncoat or *Q* is the turncoat.
 b. [From Zack]: It must be that either *Q* is the turncoat or *R* is the turncoat.

The Boss gets the messages, concluding that *Q* is the turncoat.

10. The scenario we use here is based on one used by Edgington (1995).

But is there room for this in the CIA theory? Not as things stand. Let a be the point of assessment The Boss occupies. He has no information at a about who the turncoat is, and hence the modal base $\llbracket B \rrbracket^{c,i,a}$ will contain P -worlds, Q -worlds, and R -worlds. But then it is immediate that CIA theories predict that Jack's and Zack's reports are both false at a : $\llbracket \text{must}(B)(P \text{ or } Q) \rrbracket^{c,i,a} = 0$ because $\llbracket B \rrbracket^{c,i,a}$ has some R -worlds in it; $\llbracket \text{must}(B)(Q \text{ or } R) \rrbracket^{c,i,a} = 0$ because $\llbracket B \rrbracket^{c,i,a}$ has some P -worlds in it. It is a bad idea for The Boss to conclude something he thinks is false on the basis of reports he thinks are false.¹¹

11. That is not to say that the CIA has nothing to say about these cases. They might, for example, argue that The Boss can do some metalinguistic reasoning about Jack's and Zack's truth predicates to arrive at the proper conclusion. Our point is that the CIA has to posit some additional mechanism to do the relevant work here and say why that mechanism is operative here (where we see natural information uptake instead of disagreement) and why that mechanism is not operative in the CIA's motivating cases (where we, purportedly at least, see disagreement instead of uptake).

Or perhaps—as has been suggested to us by (name withheld because the informant is under witness protection)—the CIA wants to claim that although in principle it is possible to have points of assessment occupied by individual judges, the official doctrine has always had it that *groups* of judges occupy points of assessment and that those groups must always include the speaker. Then—so they might say—this argument about uptake only applies to a very limited and uninteresting CIA theory. Allow us to wax conspiratorial: this is a cover-up. First: it is plainly false that CIA agents have insisted on any such thing. Second: it is a good thing, too. For suppose that the group of judges at a point of assessment must always include the speaker, and assume we aggregate the information available at the point in a standard way. That may help with our example (22)—since Jack knows that either P or Q is the turncoat, then at no point of assessment that includes him will his note be false, and similarly for Zack, and hence for what The Boss concludes. But the cover-up just exchanges trouble for trouble. Here is why. The classic cases of disagreement that are meant to be grist for the CIA mill go like this: A says *It might be that P* ; B disagrees, saying *It can't be that P* . If they disagree in the first place, then they must also disagree were B to speak her mind first. And, assuming that *might* and *must* are duals, that means the following should constitute a disagreement:

- (iii) a. B : It must be that not- P .
- b. A : It might be that P .

We are supposed to have a case where A and B both speak truly but disagree. But the cover-up says that the speaker of a modal claim is always one of the judges at any point of assessment for that claim. Let's say B 's claim is true at the point of assessment occupied by B alone. Since there are no worlds compatible with what B knows at which P , then for any group to which B belongs, there are no worlds compatible with what that group knows at which P . And since the point of assessment that also includes A is one such group, B 's claim is true there. That means that A has no business denying what B says. Hence the cover-up merely exchanges the difficulties in agreeing with

Next, we will consider cases in which the assessor knows more than the speaker. By the CIA's lights, an assessor j_a should reject any *might*-claim whose truth (with respect to a) would require worlds incompatible with the facts at a . But in fact, we can easily agree to an epistemic modal *might* φ in conversation even when we know whether φ . We saw an example of this already in the game of Mastermind being played between Mordecai and Pascal. After some rounds where Mordecai gives Pascal hints about the solution, Pascal says:

(23) There might be two reds.

Mordecai—who knows the solution—can reply by agreeing with the modal claim:

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

If disagreement cases are supposed to push us toward CIA theories, then this agreement should push us away. Mordecai knows exactly how many reds there are, so (23) is either false at the point of assessment he occupies (if there aren't two reds) or his reply (24) violates the maxim of quantity (if there are). But his reply is marked in no way.

7. *might* and *and/or*

Sally does not know where Joe is, but she knows he is in either Boston or New York. She says

(25) Joe might be in Boston or he might be in New York.

This entails two *might*-claims:

- (26) a. Joe might be in Boston.
 b. Joe might be in New York.

Why a disjunction of two existential modal claims entails each of its disjuncts is a bit of a puzzle, but *that* it is an entailment is clear.¹²

The point we want to make is that CIA theories predict Sally to be met with disagreement if she utters (25) in the company of anyone

must-claims for difficulties in disagreeing with them. So pick your poison. And third: relevant groups of knowers? Relevant according to what? The context of utterance has something to say here, not the point of assessment. But then, echoing a point by now familiar, we have begun the process of saving the CIA by turning it into a kludgy version of the canon.

12. See, for example, Kamp 1974; Zimmermann 2000; Simons 2005; Geurts 2005.

who knows where Joe is. The reason is that (25) entails both (26a) and (26b), and the CIA theories predict that one or the other of them will be false at any point of assessment where it is known where Joe is.

But this prediction does not square with the facts. George's reply here is absurd:

- (27) a. Sally: Joe might be in Boston or he might be in New York.
- b. George: ??Nah/That's false. He's in New York.

Just as plainly, once Sally learns that Joe is in New York, she is not going to retract her earlier claim:

- (28) a. Sally: Joe might be in Boston or he might be in New York.
- b. George: He can't be in Boston. He's in New York this week.
- c. Sally: ??Oh, then I guess I was wrong.

And eavesdroppers fare no better. Suppose that Blofeld finds the misleading evidence pointing to Bond's being in Zürich. But Blofeld is no dummy. He knows that Bond might have planted it. He can say (29a) to Number 2, but Leiter (eavesdropping with Bond from London) cannot complain with (29b):

- (29) a. Blofeld: Bond might be in Zürich or he might be in London.
- b. Leiter: ??That's false/Blofeld's wrong!

Leiter cannot disagree with Blofeld even though (29a) entails something that—given the CIA theory—is false when assessed by Leiter.

Our point is *not* that CIA theories fail to predict that (25) entails (26a) and (26b). The trouble is that, given this entailment, CIA theories straightforwardly predict that disjunctions of *might*-claims like (25) are *false* at any point of assessment where it is known where Joe is. That is a hard pill to swallow.

Or consider a quantificational variant. There has been a heist, and in the early stages of our investigation we narrow the field of suspects down to a dozen. But it is still early on, so we report to the press:

- (30) There are many people who might be the culprit.

Later, when we have eliminated all but two suspects, no one will accuse us of having said something false in the early going.

Pretty much the same point can be made by considering conjunctions of *might*-claims (though the *or*-variants tend to be more robust). Chris is hiding from Alex and Billy (there are dishes to be done):

- (31) a. Alex: She might be hiding upstairs, and she might
be hiding downstairs.
b. Billy: Right—let’s split up to find her.
c. Chris: ??Alex is wrong! I’m hiding downstairs!

Chris can’t reply—*sotto voce* or otherwise—like this. And when they find her, Alex will not be taking back what she said earlier.¹³

So far we have been pressing the point that CIA theories make implausible predictions about the semantics of disjunctions, conjunctions, and certain quantificational constructions that embed *might*-claims. There is a nearby pragmatic worry, too. It is bad conversational practice to say things that you know your conversational partners think are false. Suppose, as before, that Sally knows that Joe is either in Boston or New York but doesn’t know which it is. And suppose Sally knows that George has definite views on Joe’s whereabouts. If CIA theories were on the right track, then it would be infelicitous for her to utter (25). But it is not. Indeed, she may well use it as a premise in an argument:

- (32) Look, I know you think you know where Joe is. But we can agree on this much: Joe might be in Boston or he might be in New York. So, if he isn’t in New York, he must be in Boston.

There is some irony in the fact that cases of information asymmetry between speaker and judge cause so much trouble for the CIA since it is precisely such scenarios that are meant to be in their wheelhouse.

13. Or consider a variant of Teller’s (1972) doting grandmother:

- (iv) It might be a boy, and it might be a girl. Should I buy blue or pink?

Grandma can speak truly even if her audience knows whether it’s a boy or a girl—she may well say the conjunction of *mights* as a plea for information. And the proper response is not to tell her that she is wrong, but what color to buy. (We remain agnostic about which subset of the four possible color-gender combinations Grandma has in mind as appropriate.) But that is not the prediction that CIA theories deliver. Hawthorne (2004, 27n68) makes a similar observation.

8. Presupposition Failure Failures

Here we look at another kind of embedding: *might*-claims under factives. Factives—like *realize*—presuppose the truth of their complements. What we will see is that *might* can embed under *realize* even when both speaker and assessor think the complement of the *might*-claim is false. The CIA's prediction is that we should have presupposition failure in these cases (or, failing that, accommodation). But that is not at all what we see.

It is pretty much standard fare that presuppositions triggered in conditional antecedents project to the entire conditional construction. Both conditionals in (33) presuppose that there is no more ice cream:

- (33) a. If Sophie realizes there is no more ice cream, there will be trouble.
- b. If Sophie doesn't realize there is no more ice cream, there shouldn't be any trouble.

The factive construction, either under the negation or not, presupposes the truth of the complement. All this happens in the antecedent, but the whole conditionals carry the presupposition.¹⁴

Meanwhile, Blofeld and Number 2 are at SPECTRE headquarters plotting Bond's demise. Bond planted a bug and some misleading evidence pointing to his being in Zürich and slipped out. Now he and Leiter are listening in from London. As they listen, Leiter is getting a bit worried: Blofeld hasn't yet found the misleading evidence that points to Bond's being in Zürich. Leiter turns to Bond and says:

- (34) If Blofeld realizes you might be in Zürich, you can breathe easy—he'll send his henchman to Zürich to find you.

And he might continue:

- (35) If he doesn't realize soon that you might be in Zürich, we better get you out of here.

Bond shuts his briefcase, straightens his tie, and tells Miss Money Penny his martini may have to wait. But what he *does not* and *cannot* do is complain to Leiter that he isn't in Zürich.

This would be unexpected were the CIA agents on the right track. Bond knows he's not in Zürich. Indeed, it is common knowledge

14. See Beaver 1997 for an overview of the phenomenon of presupposition and the space of theories for accounting for it.

between Leiter and Bond that he isn't. Any plausible way of drawing the boundaries locates Bond—and Leiter—in a point of assessment *a* determining a modal base $\llbracket B \rrbracket^{c,i,a}$ that has only not-Zürich-worlds in it. We *should*, given the CIA, have about as uncontroversial a case as you can imagine of presupposition failure. But in fact we detect nothing amiss.¹⁵

Once we see this phenomenon with conditional antecedents, we can spot it in related constructions like *when* and *once* and *whether*:

- (36) Whether or not Blofeld realizes you might be in Zürich, it was excellent work planting the bug and the evidence without being noticed.

(Even Bond likes a pat on the back every now and then.)

- (37) a. Once Blofeld realizes you might be in Zürich, you can breathe easy—he'll send his henchman to Zürich to find you.
 b. When Blofeld realizes you might be in Zürich, you can breathe easy—he'll send his henchman to Zürich to find you.
 c. As soon as Blofeld realizes you might be in Zürich, you can breathe easy—he'll send his henchman to Zürich to find you.

These are far more apt to cause tie loosening and martini drinking than conversational bewilderment.

9. A Sketch of an Alternative

We have spent all of this essay laying down our gauntlet for CIA theories to run. And that is as it should be. We will leave a careful exposition of our own much more conservative treatment of the puzzles of epistemic modality to another more leisurely occasion.¹⁶ But here is a quick preview of what we have in mind.

First, though, here's what won't quite work. Within the standard theory, there is quite a bit of latitude in what might be counted as "what is known." In particular, there is agreement that at least some occurrences of epistemic *might* are sensitive not just to what the speaker knows

15. This kind of example is also trouble for other stories about epistemic modals—for example, the way of cashing out *might* as advice not to overlook the possibility that the prejacent is true that Swanson (2006) develops.

16. In particular in von Fintel and Gillies, forthcoming.

(as would be expected under a solipsistic interpretation) but to what a relevant group (usually containing the speaker) knows.¹⁷ Why not then say that the retraction/disagreement cases that CIA theories highlight are simply cases of group-based readings? We retract or disagree with a *might* claim because it is in fact *false* that the prejacent is compatible with what the relevant group knows.

There are issues with such an analysis. But not the issues you might think. The CIA agents point out that eavesdroppers can disagree with a *might* claim, even though it might be hard to argue that they are part of the relevant group about whose knowledge the speaker is making a claim. And, in fact, that is precisely why CIA agents resort to eavesdropper cases—they want to marshal some disagreement data in a family of cases that don't look ripe for explanation by appeal to a wider group of relevant knowers. But this is a pretty weak argument since in the relevant cases one *can* in fact argue that the eavesdroppers are—unbeknownst to the speaker—part of the relevant group.

But wouldn't a group-based story about epistemic modals make it pretty tough to assert them? Unless the speaker knows everything that the relevant group knows, there is a question of how the speaker can ever be justified in making a group-based *might* claim, and so how the speaker could ever live up to the norms of assertion. And if the speaker *does* know what is known to the relevant group, thereby licensing the assertion, then there would be no reason to resort to a group-based interpretation in the first place. That is a bit of a puzzle, but there is a loophole left open: one could tell a story that makes the utterance of *might* claims not subject to the usual norms of assertion, perhaps because uttering a *might* claim is not in fact an assertion but a different kind of speech act, such as a conjecture or a "recommendation not to ignore a certain possibility" (see Swanson 2006). We do not think that this is a hopeless endeavor—in fact, we think something in this ballpark needs saying since there are group-based readings of epistemic modals, something to which all parties are (or should be) agreed.

These worries don't sink a simple group-based analysis. The main problem we have with a simple group-based analysis is that it does not explain the variability of judgments that we find. Yes, speakers retract some ill-advised or doomed *might* claims, and they sometimes retract *might* claims they seemed perfectly in the right in making, but they can also stick to their guns even in the face of new evidence.

17. For relevant cases see Hacking 1967; Teller 1972; DeRose 1991.

So, what we want to develop is an analysis that centers around the variability of interpretations of epistemic modal claims. The starting point is the observation that all the trouble arises with *simple, unmodified might/must* claims. No speaker will feel pressured to retract a statement like

- (38) As far as I know at the moment, George might be
in Boston.

because she has been quite explicit in limiting her claim to the evidence available *to her at the time of utterance*. The insight we want to start from is that when a speaker does not choose to restrict her epistemic modal claim in such an explicit manner, he or she is choosing not to constrain the interpretation of her sentence. This means that her sentence could be interpreted in a variety of ways: as a solipsistic claim about her current evidence, as a claim about what a group to which she belongs currently knows, or even as a claim about all the evidence available to (but not necessarily already processed by) that group. If the context does not make it clear which interpretation is intended, what we have is contextual vagueness—it isn't clear what precisely the value of the modal base is. We would like to offer the metaphor of a “cloud of admissible contexts” with respect to which the sentence might be interpreted.¹⁸

Now, there is an interesting dynamic that is asymmetric between speaker and hearer. The hearer may be tempted to extract as much information as possible from what the speaker said, in which case the hearer will tend to interpret the sentence relative to an admissible context that makes the sentence as strong as possible. In other words, the hearer may tend to interpret a *might* claim as a statement about what evidence the relevant group has. This of course makes it easier to disagree with and reject the sentence. The speaker, on the other hand, can always retreat and say that what she intended was the narrow solipsistic interpretation (after all, she gave no explicit indication otherwise). As we saw, speakers can indeed do that and insist that new evidence doesn't invalidate their original claim.

But speakers can also choose not to fight that battle. After all, many times it simply doesn't matter whether or not they were right, whether they correctly reported on their evidential situation or not.

18. The germ of our proposal comes from a remark made by Angelika Kratzer to Kai von Fintel at a University of Massachusetts linguistic colloquium in December 2003.

What matters is whether the prejacent can still turn out to be true, and if new evidence shows that it can't, there may be no point in insisting that the original claim was true. This is where epistemic modals differ from otherwise similar facts about quantifiers. Consider:

- (39) 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.

While the retreat in (39) seems fairly natural, *might*-claims only sometimes feel worth defending via such a retreat.

We would like to point out that this “cloud of contexts” picture might also give us a handle on the norms of asserting *might* claims. The idea would be that when one asserts a sentence that is open to a variety of admissible interpretations, one needs to be justified in asserting it under the weakest admissible interpretation. In the case of *might* claims, that would typically mean that one has to be justified in asserting them as claims about one's own state of evidence. Of course, the hearer might interpret the sentence in a stronger fashion—and the speaker may well have anticipated and intended that—but the speaker can always retreat to the weakest interpretation, so that interpretation is the one she needs to be able to defend.

The main idea, then, is that the kinds of cases that CIA agents like to talk about are in fact cases where speaker and hearer adopt a strong, group-based interpretation. In general, unmodified epistemic modal claims are variable in their interpretation, which gives rise to a variety of pragmatic subtleties. So, in the relevant cases, we see a pretty expansive range of acceptable responses, ranging from retraction to sticking-to-guns. We leave it as an exercise for the reader to see how our analysis would deal with the list of cases we brought up in this essay that we argued are problematic for the CIA theories.

References

- Beaver, David I. 1997. “Presupposition.” In *Handbook of Logic and Language*, ed. J. van Benthem and A. ter Meulen, 939–1008. Cambridge, MA: MIT Press/Elsevier.
- DeRose, Keith. 1991. “Epistemic Possibilities.” *Philosophical Review* 100: 581–605.
- Edgington, Dorothy. 1995. “On Conditionals.” *Mind* 104: 235–329.

- Egan, Andy. 2007. "Epistemic Modals, Relativism, and Assertion." *Philosophical Studies* 133: 1–22.
- Egan, Andy, John Hawthorne, and Brian Weatherson. 2005. "Epistemic Modals in Context." In *Contextualism in Philosophy*, ed. G. Preyer and G. Peter, 131–69. New York: Oxford University Press.
- von Fintel, Kai, and Anthony S. Gillies. Forthcoming. "Might Made Right."
- Geurts, Bart. 2005. "Entertaining Alternatives: Disjunctions as Modals." *Natural Language Semantics* 13: 383–410.
- Gibbard, Allan. 1981. "Two Recent Theories of Conditionals." In *Ifs*, ed. W. L. Harper, R. Stalnaker, and G. Pearce, 211–47. Dordrecht: D. Reidel.
- Hacking, Ian. 1967. "Possibility." *Philosophical Review* 76: 143–68.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kamp, Hans. 1974. "Free Choice Permission." *Proceedings of the Aristotelian Society* 74: 57–74.
- Kratzer, Angelika. 1981. "The Notional Category of Modality." In *Words, Worlds, and Contexts: New Approaches to Word Semantics*, ed. H. J. Eikmeyer and H. Rieser, 38–74. Berlin: de Gruyter. Reprinted in *Formal Semantics: The Essential Readings*, ed. P. Portner and B. H. Partee, 2002. Oxford: Blackwell.
- . 1986. "Conditionals." *Chicago Linguistics Society* 22: 1–15.
- Laserson, Peter. 2005. "Context Dependence, Disagreement, and Predicates of Personal Taste." *Linguistics and Philosophy* 28: 643–86.
- MacFarlane, John. 2004. "Epistemic Modalities and Relative Truth." <http://socrates.berkeley.edu/~jmacf/epistmod.pdf>.
- Simons, Mandy. 2005. "Dividing Things Up: The Semantics of *Or* and the Modal/*Or* Interaction." *Natural Language Semantics* 13: 271–316.
- Stephenson, Tamina. Forthcoming. "Judge Dependence, Epistemic Modals, and Predicates of Personal Taste." *Linguistics and Philosophy*.
- Swanson, Eric. 2006. "Something *Might* Might Mean." Unpublished manuscript, University of Michigan.
- Teller, Paul. 1972. "Epistemic Possibility." *Philosophia* 2: 203–320.
- Zimmermann, Thomas Ede. 2000. "Free Choice Disjunction and Epistemic Possibility." *Natural Language Semantics* 8: 255–90.