

Craige Roberts
May, 2015

The Character of Epistemic Modality: Evidentiality, indexicality, and what's at issue¹

ABSTRACT: I argue that statements containing English epistemic modals are sensitive to context in three distinct ways, crucial to understanding the complex patterns of interpretation and response attested in the literature. First, I endorse an evidential Kratzerian semantics, as originally proposed by von Stechow & Gillies (2010). But I modify their semantics so that *must* and *may* are merely doxastic and suppositional instead of truly epistemic, yielding a truth conditional semantics from which one of their key stipulations is argued to follow. Second, I argue that *must* and other epistemic modals are indexically anchored to a discourse center—an agent-at-a-time whose doxastic state is currently under discussion in the context of utterance (Roberts 2015). This correctly predicts the contextually limited range of candidates for the doxastic anchoring agent of such a modal, as attested in the literature, thereby constrains what body of evidence is understood to be relevant, and helps to explain the modal's apparent scope in a given context—a pseudo-scope effect rather than a reflection of syntactic scope at LF (*pace* Hacquard 2013). Third, in most contexts, where the epistemic state of the anchoring agent is not itself at-issue in the sense of being RELEVANT to the Question Under Discussion (Roberts 1996/2012), the evidential content associated with the modal will not be the target of an apt response. Such not-at-issue uses parallel parenthetical uses of attitude reports (Simons 2006). Only when the evidentiality itself *is* at-issue, as in von Stechow & Gillies' (2007) Mastermind contexts, does a direct response target that content. This account explains a variety of other properties of epistemic modals and sheds light on some outstanding puzzles, including Yalcin's (2007) version of Moore's paradox for embedded epistemic modals, and purported arguments for modal Relativism (Egan, Hawthorne & Weatherson 2005). And it argues for a more refined notion of RELEVANCE, which considers not only truth, but the possibility and probability of potential answers.

1. Introduction

Some of the most interesting studies on the relationship between semantics and pragmatics involve simple function words—members of closed classes, with a very high frequency of

¹ I am indebted to the members of the OSU project on Perspectival Expressions—Jefferson Barlew, Greg Kierstead, and Eric Snyder—for discussions of this material over many months and their own stimulating exploration of related ideas. And to David Beaver, Mandy Simons and Judith Tonhauser for discussions over many years about the notion of at-issueness and its role in interpretation, as reflected in the analysis in §4.5 of this paper. Some of the material in §4.5.2 was included in a joint presentation with Greg Kierstead at QiD Frankfurt in 2014 (Kierstead & Roberts 2014). I am also very grateful to audiences at the Rutgers University ErnieFest in 2014, the MASZAT group at the Research Institute for Linguistics of the Hungarian Institute of Sciences in spring 2015, and at ReDRAW '15 at the University of Groningen; and to Kai von Stechow, Hans-Martin Gärdner, Thoni Gillies, Jack Hoeksma, Ernie LePore, Emar Maier, Friederike Moltmann, and Jennifer Spender for stimulating discussions and comments. The perspective project received invaluable support in 2013-14 from a Targeted Investment in Excellence grant from OSU, a Research Enhancement Grant from the OSU Colleges of the Arts and Sciences, and a supplement to NSF Grant #0952571, the latter originally awarded to Beaver, Roberts, Simons & Tonhauser. This manuscript was completed while I was a Senior Fellow in 2014-15 at the Institute for Advanced Studies at Central European University, Budapest, Hungary, sponsored by Budapesti Közép-Európai Egyetem Alaptvány, and again, I am deeply grateful for their support, and for the assistance of OSU, without which I could not have accepted the fellowship. The theses promoted herein are the author's own, and do not necessarily reflect the opinion of the CEU IAS or any of the other sponsoring organizations or individuals cited here.

occurrence in the language. Pronouns and definite articles, connectives like *or* and its translation counterparts across languages, *only* and its kin, and many others are all clearly context-sensitive in their semantics, and in fact it is arguably this constrained context-sensitivity that lends them their power and flexibility. About the semantics and pragmatics of each there is a rich literature but as yet no consensus. Here I offer a study of two members of another closed class of context-sensitive items, the English epistemic modal auxiliaries, focusing here on *must* and *might*.

I make no pretense of having the last word on either, let alone on the full range of modal auxiliaries and other epistemic vocabulary. But drawing on a rich existing literature full of insights and puzzles, I offer some new observations and integrate them into a semantics for these English exemplars in which they are context-sensitive in ways that help to explain these puzzles. In keeping with the general Kratzerian approach to natural language modal auxiliaries, their proffered content is very simple—effectively universal (*must*) or existential (*might*) quantification over a relevant set of possible worlds. The interesting work is done by other aspects of their conventional Characters: what is presupposed, that is to say, how an epistemic modal conventionally appeals to the context of utterance to retrieve crucial features of its meaning in that context. Again, the presupposed content of an epistemic modal auxiliary (EMA) is itself rather simple, but because discourse is rich and complex, so that different contexts can vary along many distinct parameters, the resulting patterns of interpretation accordingly show complex patterns of variation.

As a first sketch, here is the proposed semantic Character of *must*:

The Character of English epistemic *must*: [informal]

Presupposed content:

Indexical anchor: There is a particular doxastic agent *a*, whose relevant belief state at a given time *t* the speaker intends to anchor the interpretation of the modal.

Evidential Modal Base: The domain of the modal operator associated with *must* is determined by a consistent set of suppositions that is a proper superset of *a*'s beliefs at *t*.

Proffered content: The modal's complement, its prejacent, is true in all the worlds given by its modal base.

Beliefs and evidence are believed by and evident to particular agents. So the use of an EMA presupposes that its interpretation is anchored by a particular, contextually salient doxastic agent. The default for an assertion is the speaker, whose beliefs are always under consideration in discourse; or for a question, the addressee, the individual whose views are being solicited. But the anchor might be 'we', the interlocutors' Common Ground (CG) being their joint *purported* belief state in the discourse and as such always (ideally) evident and salient. Or when the EMA occurs in the complement of an attitude predicate, the anchor will most likely be the agent of the attitude. As we will see, there is independent evidence that at any given time in discourse there are a limited set of possible doxastic anchors; Roberts (2015) argues that these generally anchor the interpretation of indexicals, including *inter alia* indexical pronouns and adverbs, and predicates like *come* (Barlew 2015). Hence, the Character of *must* is indexical. Given the general nature, behavior, and availability of indexical anchors, this predicts both contextual constraints on *and* clues to the intended resolution of the indexical anchoring presupposition.

Standard Kratzerian semantics for natural language modality takes such modals to be interpreted relative to an often-implicit modal base MB, restricting the domain of the modal. It is sometimes assumed that the MB for a given modal is freely chosen. But it seems that particular modal auxiliaries carry a variety of lexical restrictions on the possible types of MB available for their interpretation. E.g., *needn't* (a modal Negative Polarity Item) can only be deontic—it cannot have an epistemic interpretation; *will/would* cannot have a deontic interpretation; *might* doesn't seem to have a dynamic interpretation; etc. (Palmer 1990). One might regard these constraints as a species of lexical subcategorization restrictions. Here I will model them as presuppositions of any speaker who uses the modal: for the EMAs under consideration here, the presupposition of an evidential modal base. This is, of course, quite similar to, and in fact inspired by, the evidentiality of EMAs argued for by von Fintel & Gillies (2010). But the version here differs from their proposal in important ways, which will be clearer when we consider the formal proposal, below.

I take proffered content (Roberts 1996) to be that aspect of the conventionally given Character of an expression which plays a role in the compositional, truth conditional interpretation of any utterance in which it occurs—what is asserted as the canonical use of a declarative, what's questioned in an interrogative, what's suggested by an imperative. The proffered content of an EMA is quite simple: the appropriate modal operator—for *must*, necessity—takes as its argument the prejacent. In keeping with Kratzer's approach, this is what epistemic *must* has in common with its deontic interpretation, as well as with *should*, *ought*, *will* and *would*—all have the force of necessity, taking the prejacent as semantic complement. Thus, I take it that the important differences between natural language modals and their logical counterparts lies in their presupposed content, and in how their presuppositions are resolved in context.

The proffered content of an EMA includes both the evidential claim associated with the operator and what the evidence bears on, the prejacent. In some contexts, responses to an EMA assertion seem to target the prejacent alone, commenting on its truth or falsity.

(1) [Context: One prosecutor talking with another about their case:]

A: Given the evidence I've seen, the victim might have known the killer at Yale.

B: No, he didn't. They weren't there at the same time.

Here, B's response seems to be a denial of the prejacent *the victim knew the killer at Yale*, and not merely of its possibility relative to A's available evidence. This has led some to argue that the EMA may act as a speech act modifier in such assertions or that there are two speech acts in such assertions or that the EMA has a special performative use reflected here (Lyons 1977, von Fintel 2003, Swanson 2006, von Fintel & Gillies 2007a, 2007b, Portner 2007a). I argue that we can instead explain the attested response patterns in terms of what's at-issue in the context of utterance (Amaral, Roberts & Smith 2007; Roberts et al. 2009; Simons et al. 2011; Tonhauser et al. 2012, Simons et al. to appear). The at-issue content of an utterance is that part of its content which is intended to address the *Question Under Discussion* (QUD) (Roberts 1996/2012). In a typical context of utterance, only the prejacent of an EMA assertion is at-issue. But in certain contexts, illustrated by von Fintel & Gillies' (2007a) Mastermind examples, the evidentiality itself is at-issue. Then in those contexts, apt response targets the EMA as well. Thus, EMAs are an especially interesting illustration of how proffered content can shift between foreground and

background in context, how this is reflected in our sense of what is asserted, and accordingly, in apt response patterns. This feature of the proposal is closely related to arguments in Moss (2015) about the interpretation of nested epistemic modals, though she doesn't focus on response patterns.

Summarizing, English epistemic modal auxiliaries *must* and *might* are context-sensitive in three ways:

- a) An EMA presupposes a doxastic anchor for its evidentiality, which must be one of a limited set of contextually available discourse centers.
- b) An EMA's modality is evidential, the nature of the evidence contextually implicated and presupposed to be available to its doxastic anchor.
- c) The proffered implications of an EMA may shift between background and foreground, as a function of the QUD.

These three parameters are relatively independent of each other, though their resolution in a particular context of utterance may be interdependent. But together, I will argue, they serve to explain a number of outstanding puzzles in the literature, summarized in the following table. In column (b), the symbol © represents a discourse center, an agent at a time whose doxastic state is under consideration in the discourse at a given point:

Contextual parameters: Puzzles (with § numbers):	(a) Evidentiality	(b) Anchoring to ©	(c) Shifting background/foreground
4.1.1 Weak necessity	√		
4.2. Variable anchoring		√	
4.3.1 Puzzling inferences		√	
4.3.2 Yalcin's puzzle	√	√	
4.4. Apparent scope		√	
4.5.1 Faultless disagreement		√	√
4.5.2 Response patterns			√

Table 1: Explaining puzzles about epistemic modals

The problem of weak necessity (§4.1) has motivated a great deal of discussion in the literature on *must*, summarized in von Fintel & Gillies (2010). As in their work, the evidentiality adopted here is intended to explain the relevant observations while retaining a basically strong semantics for *must* and comparable semantics for *might*. Similarly, many have puzzled over the apparent variability in anchoring of EMAs (§4.2), including Hacking (1967), DeRose (1991), Egan et al. (2005), MacFarlane (2006), von Fintel & Gillies (2007a), and Yalcin (2007). Anchoring to a discourse center ©, given independently motivated constraints on the availability of discourse centers, is intended both to explain and predict variability, and to constrain it appropriately. I will also consider interactions between EMAs and attitudes (§4.3), and consider how these bear on recent arguments for a non-standard semantics for EMAs, focusing on arguments in Egan et al. (2007) (§4.3.1) and Yalcin (2007) (§4.3.2), and showing how the presupposed doxastic anchoring and associated evidential semantics proposed for *must* and *might* can explain the puzzles these authors consider while maintaining a simple, Kratzerian semantics. Some have claimed that EMAs inevitably take wide scope (e.g., Hacquard 2006, 2013), and Hacquard (2013)

provides good evidence that this tendency is cross-linguistic. But I argue (§4.4) that (a) that is not always the case, since EMAs can take narrow scope relative to a wide range of operators, and (b) we can explain the tendency to wide scope through a combination of the indexicality of EMAs and the phenomena of pseudo-scope, wherein operators which are part of anaphoric expressions have their apparent scope fixed as a function of the antecedent resolving the anaphora. Finally, I consider examples which display evidence of faultless disagreement and “faulty agreement” (§4.5.1) and those displaying a variety of patterns of felicitous response to statements with EMAs (§4.5.2), responses which sometime seem to target the prejacent alone, and other times the entire evidential claim. I provide an alternative characterization about what it is to agree about EMA statements, and evidence that the pattern of felicitous response in a given example is a function of what’s at-issue in the context of utterance, undercutting a variety of criticisms of the standard semantics which trade on these phenomena.

In what follows, in §2 I give some necessary background for the semantics to be developed, explaining my assumptions about the nature of a context of utterance, and especially about the crucial new notion of a discourse center ©. In §3, I offer and explain formal Characters for *must* and *might*. In §4 I illustrate the implications of these Characters for utterances including EMAs, focusing on the puzzles from the literature just outlined. Finally, in §5, I briefly present a few conclusions. Throughout it should be kept in mind that this proposal is not aimed at resolving all features of the semantics of epistemic modals. For example, it does not address what Moss (2015) calls “graded epistemic modals” like *probably* and *likely*, and how these interact with *must* and *might*. But I take it that not only must any adequate semantics for the EMAs take the three kinds of context-sensitivity noted above into account, but that accounts which do so can be significantly improved and simplified in other respects.

2. Background: Perspectival content and doxastic centers

In any given discourse at any given time, there is a limited set of points of view that are relevant and salient at that time. These points of view are doxastic, they consist of the evident (purported) beliefs those of an agent *a* at a time *t*.² Because one of the functions of an agent’s point of view is to play a role in characterizing *de se* attitudes, we take such an agent-at-a-time, $\langle a, t \rangle$ to constitute a center in the technical sense of Lewis (1979). Roberts (2015) argues that interlocutors track the set of entities whose point of view is relevant at a given time, the set of **discourse centers** at that point in the discourse. So a salient doxastic center is not only an entity-at-a-time in a world, but is reflected in the special status of a corresponding discourse referent. To identify a discourse center ©, we give it the two indices corresponding to those on the discourse referents for the agent and time, both of which must be in the set of familiar discourse referents in order for © to be a discourse center; $\textcircled{c}_{i,j} = \langle d_i, t_j \rangle$. The doxastic state of (the denotation of) a center $\textcircled{c}_{i,j}$ is the set of centered propositions consistent with the beliefs of (the agent which is the value of) d_i at (the time which is the value of) t_j . Those available in a given context always include the speaker at utterance time: \textcircled{c}^* , and the addressee at utterance time: $\textcircled{c}^@$. And it includes their join, inclusive ‘we’: $\textcircled{c}^* \oplus \textcircled{c}^@$, reflecting the point of view captured in

² As Stalnaker and others have been at great pains to explain, these are the beliefs that the interlocutors in the conversation take the relevant agent to hold. But they may be wrong, even misled, about that agent’s belief state. For the purposes of the conversation, that doesn’t matter.

the Common Ground CG. But under the scope of an attitude predicate P , the doxastic state of the agent of the attitude also becomes relevant: \odot^P , while in Free Indirect Discourse (FID),³ the doxastic state of the agent whose viewpoint at the time of the reported attitude is adopted serves as an additional center \odot^{FID} . The set of available doxastic centers in a given discourse at a given time is \odot_D , the discourse centers at that time in D , which, the Reference Time in Partee (1984), is updated regularly over time as the speaker and addressees change, we move in and out of the scope of attitudes or FID, etc. Summarizing, here's some terminology:

A **doxastic point of view**: the doxastic attitude of an agent A at a time in a given situation (or world)

A 's **doxastic perspective**: her associated doxastic state—a set of centered worlds reflecting what A believes at that time in the world of evaluation. As in Lewis (1979), the centers in these worlds are used to capture what Stalnaker (2008) calls “self-location in thought”.

A at the time she holds that attitude is the **center** of that doxastic state.

In any discourse, the **speaker** Sp at the time of utterance t^* plays the role of a distinguished discourse center, $\odot^* = \langle Sp, t^* \rangle$, intentionally reporting on her purported doxastic state. The **addressee** Add at utterance time corresponds to another, $\odot^@ = \langle Add, t^* \rangle$. The Context Set corresponding to their Common Ground (a joint purported belief state, see Ninan 2010) corresponds to yet another, the inclusive ‘we’: $\odot^{\text{CG}} = \langle Sp \oplus Add, t^* \rangle$.

But **other centers** may be introduced, either conventionally—such as the explicit agent of a self-locating attitude predicate like *believe*: \odot^{believe} —or pragmatically—as in Free Indirect Discourse: \odot^{FID} .

To be an **indexical** expression is to have a Character which conventionally presupposes a contextually given doxastic center as **anchor**. Some indexicals lexically constrain their admissible anchors, typically to \odot^* or $\odot^@$, or perhaps as well to (as in Amharic, Schlenker 2003) the agent of a verb of speaking. Anchoring is a type of anaphora, but the presupposed anchor typically isn't coreferential with the indexical expression's proffered content.

To define these notions more precisely, we first need to define the notion of a doxastic state, and the associated notion of a doxastic accessibility relation over possible worlds. Stalnaker (2008) argues for a modification of Lewis' characterization of centered worlds, developed to account for *de se* phenomena, and I adopt Stalnaker's version here without modification. He has multiple reasons for this, but the one I find most compelling is what he calls a problem of *calibration*, crucial to comparing cognitive states. Hintikka's approach to propositional attitudes via modal accessibility relations makes possible comparison of the content of the objects of such attitudes across times and across persons: Two individuals A and B (or one individual at two times) believe the same proposition p at t in w just in case both $\text{DOX}(A, t, w) \subseteq p$ and $\text{DOX}(B, t, w) \subseteq p$. Then:

The unreconstructed Hintikka-style models of cognitive states took calibration for granted, but ignored belief change and self-location. The Lewis centered-worlds models recognized self-location, but provided no resources for representing the relations between informational states across time and across persons, and so no resources for clarifying the dynamics of knowledge and belief, or the communication of information between different subjects. The general framework

³ See Eckardt (2014) for a recent formal treatment of FID.

that I am promoting allows for calibration across time, and across different subjects, but it also recognizes that calibration is a nontrivial problem, and may not be well defined in all cases. [Stalnaker 2014]

If beliefs are sets of centered worlds, and two distinct agents' beliefs involve sets with distinct centers, how can we compare what they believe? If the same agent at distinct times corresponds to two distinct individual-time pairs, two centers, how can we compare what that agent believes at different times?

Stalnaker points out that replacing worlds in the Hintikka approach with centered worlds (below) permits us to talk about just one doxastic accessibility relationship R , the agent given by the center of its first argument. Hence, a belief state involving Lewis' centered worlds can be characterized thus:

a **belief state** is a pair consisting of a centered world and its Dox-related belief set:

- the **base (centered) world**: the determining centered world, an ordered pair consisting of (a) the center: a person whose beliefs are being represented and the time at which she has those beliefs, and (b) the possible world in which the center has those beliefs
- the **belief set**: the determined set, a set of centered worlds of the same type as the base world. In each pair $\langle c, w \rangle$, the c represents what the base subject takes herself to be in w , a world which, for all she believes, may be actual. The worlds in these pairs are those which would be accessible from the base world under Hintikka's doxastic accessibility relation relativized to the base center.

Stalnaker's modified theory of centered worlds is realized with a model $\langle W, S, T, \geq, E, R \rangle$ where:

1. W is a nonempty set of possible worlds
2. S is a set of *subjects* or believers [my *doxastic agents*]
3. T is a set of times
4. \geq is a binary transitive connected anti-symmetric relation on T , a relation that determines a linear order of the times.
5. E is the set of centered worlds meeting the condition that the subject of the center exists in the world at the time of the center, where
 - A **center** is a pair, $\langle A, t \rangle$, where $A \in S$ and $t \in T$. Subjects may exist at some times at some worlds, and not at others.
 - A **centered world** is a pair $\langle c, w \rangle$, where c is a center and $w \in W$.
6. R is a binary relation on E that is transitive, Euclidean and serial. R must also satisfy condition (*), below. To say that $\langle \langle A, t \rangle, w \rangle R \langle \langle B, t^* \rangle, w' \rangle$ is to say that it is compatible with what A believes at time t in world w that she is in world w' , that she is person B , and that the time is time t^* .

R in 6 is a doxastic accessibility relation (the sort of relation I call *DOX* throughout this paper), representing a subject's beliefs at a time in a world. The requirements on R guarantee that the agent has access to what she believes and to what she does not. See Stalnaker (2014, Chapter 2) for very useful discussion of the properties of this model. It has the additional condition (*):

- (*) For any centers, c^* , c' and c'' , and worlds w and w' : if $\langle c^*, w \rangle R \langle c', w' \rangle$ and $\langle c^*, w \rangle R \langle c'', w' \rangle$, then $c' = c''$.

Intuitively, (*) tells us that “ignorance or uncertainty about where one is in the world is always also ignorance or uncertainty about what world one is in” [2012:70]. Stalnaker tells us that (*) is the main respect in which this model differs formally from Lewis’, in which the same center was permitted to occur in two “places” in the same world. Lewis argued that this was necessary in order to properly model the structure of *de se* beliefs, which he argued to be inherently more fine-grained than could be captured with possible worlds. But because of (*), for Stalnaker the contents of a belief state “can be taken to be ordinary propositions—sets of uncentered possible worlds, even though the centers determined by a particular belief state may play a role in determining which proposition is denoted by a that-clause with indexical expressions in it” (2008:71). That is, the “distinctive self-locating character [of self-locating beliefs] will be a feature of the subject’s relation to that content, and not a feature of the content itself” [my emphasis].

Summarizing, a doxastic accessibility relation *DOX* takes an agent a , a time t , and a world w , $\langle \langle a, t \rangle, w \rangle$ —the centered world whose center is $\langle a, t \rangle$, and yields a set of centered worlds, those in which every proposition that a believes at t in w is true and the center reflects a ’s self-location in that world. This is the agent’s *belief state*. And since the agent self-locates as the centers in her belief state worlds, via those centers we can characterize his *de se* beliefs.

Now we can use these notions to give technical definitions for the intuitive notions discussed above:

- A **doxastic center** is an ordered pair consisting of a doxastic agent a and a time t : $\langle a, t \rangle$.
- A **doxastic origin** is an ordered pair of a doxastic center and a world: $\langle \langle a, t \rangle, w \rangle$.
- A **doxastic point of view** is a doxastic origin and its associated doxastic relation *DOX*.
- A **doxastic perspective** is the information accessible from a doxastic point of view, a ’s belief set at t in w : $\text{DOX}(\langle \langle a, t \rangle, w \rangle)$.
- A **proposition** is a set of centered worlds. (Lewis 1979, Stalnaker 2008)

I also adopt the following notion of the context of utterance:

Context of utterance in a discourse D: $\langle \text{DG}_D, \text{QUD}_D, \text{CS}_D, \text{DR}_D, \text{©}_D \rangle$,⁴ consisting of

- DG: the interlocutors’ evident goals, their Domain Goals
- QUD: the set of questions currently under discussion
- CS: the interlocutors’ Context Set, the set of worlds compatible with their CG
- DR: the set of Discourse Referents (DRefs)⁵
- ©: the set of discourse centers, each the ordered pair of a DRef and a time: $\langle d, t \rangle$.

As in Roberts (1996/2012), the QUD plays a central role in constraining what a speaker can reasonably be taken to mean_{nn} (in the sense of Grice 1957) by her utterance. This is driven by a requirement of RELEVANCE to the QUD (see that paper for details):

⁴ In what follows, I often ignore DG and QUD for a given context when these are not relevant.

⁵ In my own recent work (Roberts 2015), I have argued that DRefs are of type $\langle s, e \rangle$. Here I assume type e for simplicity, since the issues to be addressed do not appeal to the richer type.

RELEVANCE to the QUD: Felicity of utterance requires RELEVANCE to the QUD, where

- a. An assertion is relevant to a QUD iff it contextually entails a partial or complete answer to the QUD.
- b. A question is relevant to a QUD iff it has an answer which contextually entails a partial or complete answer to the QUD.
- c. A directive is RELEVANT to a QUD iff its realization promises to play a role in resolving the QUD

The Character⁶ of a constituent, e.g. a lexical item, may involve at least three types of content:

Character consists of three types of conventional content:

Presupposed content: that which constrains the contexts of utterance in which utterance of the content is felicitous (cf. Heim 1983)

Auxiliary content:⁷ that which is directly attributed to some discourse center (Potts 2005, Anderbois et al. 2015, Martin 2014, Barlew 2015)

Proffered content: that which enters into the compositionally calculated truth conditions of the utterance in which it occurs, which in turn constrains what is asserted, asked or directed (Roberts 1996/2012)

Then we can define a notion that will play a role in differentiating background from foreground proffered content in an utterance:

at-issueness:

For any proposition p , let $?p$ denote the question whether p , i.e. the partition on the set of worlds with members p and $\neg p$. Then:

A proffered proposition p is **at-issue** relative to a question Q iff $?p$ is RELEVANT to Q .
[revised from Simons et al. 2011]

The revision amounts only to the addition of the term *proffered* in the definition, thus explicitly excluding presupposed and auxiliary content from what could be at-issue. Another way of saying this is that proffered content is that which contributes to the move being made by an utterance including the expression in the language game at play. It is only the move itself, and not constraints on its felicity or auxiliary content, which constitutes the speaker's contribution to the resolution of the QUD, i.e. to what's at-issue. Of course, this leaves open the important question of what's conventionally presupposed by an expression. My colleagues and I have argued (Roberts et al. 2009, Simons et al. 2011, Tonhauser et al. 2013, etc.) that much that's

⁶ The term is, of course, borrowed from Kaplan (1979), but the character of Character here is rather different. In particular, in keeping with dynamic theories of context generally, I assume that context may be updated in the course of interpretation, so that the presupposed content of Character needn't be satisfied by the context of utterance alone, but may be merely locally satisfied, as in Kamp & Reyle (1993), Martin (2013), AnderBois et al. (2015), *inter alia*.

⁷ Auxiliary content includes the content associated with Potts' (2005) Conventional Implicatures, though I'll have nothing to say about those here. Note that Amaral, Roberts & Smith (2007) and Harris & Potts (2009) provide evidence that those contents are anchored to a salient point of view, frequently, though not necessarily, that of the speaker or addressee (*pace* Potts 2005).

taken to be presupposed is instead Not at-issue proffered content. This comes to bear on the analysis of the EMAs to be developed below, especially in §4.5.

The set of **discourse centers** \odot_D : a dynamically changing set of the familiar doxastic centers (agents at a time) whose doxastic perspective the interlocutors take to be relevant at a given point in the discourse:

$\odot_D \subseteq \{ \langle d_i, t_j \rangle \mid d_i, t_j \in DR \ \& \ d_i \text{ is a doxastic agent whose beliefs at } t_j \text{ are under discussion in } D \}$. Further:

- \odot_D always includes a distinguished center $\odot_{i,j}^*$, corresponding to the speaker d_i at the time of utterance t_j , as well as one $\odot_{k,j}^@$ corresponding to the addressee at that time.
- other centers are introduced conventionally (under attitude predicates) or conversationally (in FID or modal subordination), in conjunction with the interlocutors' consideration of alternative doxastic states.

Roberts (2015) constitutes an extended argument for the utility of discourse centers, and of the corresponding *de se* semantics and pragmatics, in the semantics of indexical expressions, including the canonical indexicals and demonstratives that were the focus of Kaplan (1979) and so much subsequent work. The arguments for this treatment include the *de se* character of indexicals (see also Wexler 2010), the treatment of so-called “fake indexicals” (e.g., Kratzer 2009), and its utility in analyzing indexicals in languages which permit them to shift from the interlocutors. As an illustration, here are the *de se* semantics for English *I* and for the Amharic first person pronominal –*ññ*:

Character of English *I*:

Given a context $D = \langle CS_D, DR_D, \odot_D \rangle$, with $\odot^* = \langle d_k, t \rangle$ a distinguished element of \odot_D :

Presupposed content: Use of I_i is felicitous in D at time t just in case $d_i \in DR_D$ and for all CS-consistent assignments g , $g(d_i) = g(d_k)$.

Proffered content: Where felicitous, for all CS-consistent assignments g , $|I_i|^{D,g} = g(d_i)$.

Note two features of this Character. First, though *I* needn't be coindexed with the distinguished discourse center \odot^* corresponding to the speaker at utterance time, the Presupposed content guarantees that it will be coreferential with the agent of \odot^* under any assignments of values to variables that are consistent with the interlocutors' joint information. This presupposition guarantees that the denotation will always be the actual speaker. Second, because \odot^* is presupposed, the resulting semantics is different from that of the definite description *the speaker* in just the right way (I argue) to guarantee the differences under modality observed by Kaplan (1979). Unlike his Character, Characters in the present framework may appeal to contexts which are updated under the scope of operators like negation or modals. However, anchoring to \odot^* will guarantee that *I* always has the effect of widest scope.

This is different from the semantics for languages which have shifted indexicals, like Amharic (Schlenker 2003):

Character of Amharic 1st person shiftable indexical $\tilde{n}\tilde{n}$:

Given a context $D = \langle CS_D, DR_D, \odot_D \rangle$:

Presupposed content: Use of $\tilde{n}\tilde{n}_i$ is felicitous in D at time t just in case there is a $\odot = \langle d_k, t \rangle \in \odot_K$ s.t. $\odot = \odot^*$ or $\odot = \odot^{\text{say}}$, and a $d_i \in DR_D$, and for all CS-consistent assignments g , $g(d_i) = g(d_k)$.

Proffered content: Where felicitous, for all CS-consistent assignments g , $|I_i|^{D,g} = g(d_i)$.

The crucial difference between the Characters of I and $\tilde{n}\tilde{n}$ is that the latter permits the anchoring discourse center to be either \odot^* or the agent of the verb meaning ‘say’. Since such an agent is only introduced under the scope of that verb, the shifted use is only licensed under ‘say’.⁸ In other languages that have shifted indexicals, they may be licensed as anchored to a broader set of types of discourse center; see Roberts (2015) for extended discussion of the relevant literature.

The Characters of the first person indexicals make them coreferential with the agent of the presupposed discourse center. But that is not in general the case when an expression is indexically anchored to a discourse center, as we will see now in the semantics of the EMAs. See also Barlew (2015) for an account of the semantics of *come* based on this type of anchoring, which he calls *perspectival*, arguing that the notion of indexical anchoring has a broad application across classes of expressions of many types.

3. Doxastic semantics and pragmatics for evidential modals

Recall the informal characterization of the meaning of *must* in the introduction. With the tools discussed in the previous section, we can now give the formal version of its Character as follows:

Character of English epistemic *must*:

Given an utterance $must_{i,j} p$ in context $D = \langle CS_D, DR_D, \odot_D \rangle$ with world and time of evaluation w , t_j :

Presupposed content:

Indexical anchor: There is a $\odot_{i,j} = \langle d_i, t_j \rangle \in \odot_D$.

Evidential Modal Base: There is a function f mapping centered worlds to sets of centered propositions (each a set of centered worlds) s.t. $f(\langle \odot_{i,j}, w \rangle) = S$, where S is s.t.:

- | | |
|--|---|
| $\cap S \neq \emptyset$ & | ‘ S is consistent’ |
| $\forall p \in S$: suppose($p, \odot_{i,j}, w$) & | ‘all p in S are supposed by d_i at t_j in w ’ |
| $\cap S \subset \text{Dox}(\odot_{i,j})(w)$ | ‘ S properly entails d_i ’s beliefs at t_j ’ |

Proffered content: $\lambda p_{\langle s, t \rangle} \lambda \langle \odot, w \rangle. p \subseteq \cap f(\langle \langle d_i, t_j \rangle, w \rangle)$

We will call the presupposed discourse center $\odot_{i,j}$ is the **doxastic anchor** for *must*, because it is $\odot_{i,j}$ ’s belief state that constrains the presupposed evidential modal base. This center may be any of those available in the context at the time of utterance, as we’ll discuss below in §4.2. The Modal Base f will be a proper extension of the centering agent d_i ’s belief state at t_j in the world of evaluation w : The conditions on the presupposed function f tell us that for $\langle \odot_{i,j}, w \rangle$ the

⁸ I have not done fieldwork on Amharic myself, so cannot report, but it would be interesting to which whether it is possible in cases involving extended speech reports—rather like modal subordination contexts—to continue to use $\tilde{n}\tilde{n}$ in those extended reports.

function yields a set of propositions consisting of all those believed by $\odot_{i,j}$ in w plus a set of propositions S merely supposed by $\odot_{i,j}$ in w .⁹ As we will see, it is this extension of the center's belief state to include merely supposed propositions which yields the evidentiality of the EMA, so that assertion of *must* φ does not entail that the speaker is committed to the truth of φ itself.

The proffered content of *must* is very simple, as in Kratzer semantics: The prejacent must be true in the presuppositionally restricted domain.

We have a comparable semantics for epistemic *might*, differing, as usual, only in that it merely requires that the prejacent be consistent with the modal base, instead of being entailed by it:

Character of English epistemic *might*:

Given an utterance *might* _{i,j} p in context $D = \langle CS_D, DR_D, \odot_D \rangle$:

Presupposed content:

Indexical anchor: There is a $\odot_{i,j} = \langle d_i, t_j \rangle \in \odot_D$.

Evidential Modal Base: There is a function f mapping centered worlds to sets of centered propositions (each a set of centered worlds) s.t. $f(\langle \odot_{i,j}, w \rangle) = S$, where S is s.t.:

- $\cap S \neq \emptyset$ & 'S is consistent'
- $\forall p \in S$: suppose($p, \odot_{i,j}, w$) & 'all p in S are supposed by d_i at t_j in w '
- $\cap S \subset \text{Dox}(\odot_{i,j})(w)$ 'S properly entails d_i 's beliefs at t_j '

Proffered content: $\lambda p_{\langle s, t \rangle} \lambda \langle \odot, w \rangle. p \cap [\cap f(\langle \langle d_i, t_j \rangle, w \rangle)] \neq \emptyset$

What is it for an agent to *suppose* that p ?¹⁰ The following type of example argues that supposition is weaker than belief:

- (2) A: It's better to have dessert before dinner than after. [= p]
- B: Well, I suppose so.

(2B) has the flavor of a grudging admission: The speaker concedes that it's reasonable to assume p , and perhaps even that it seems likely or that the evidence suggests that it's true. But B seems to hint that she's not yet entirely convinced or committed to the truth of p .

Similarly, the following sequences seem consistent:

- (3) a. Mark supposed that he would have to go through with it.
- b. But he still hoped that there might be another way.
- or

⁹ The presupposed MB f could be characterized more generally: for *any* given $\langle \text{center}, \text{world} \rangle$, it maps that center to a set of propositions s.t. it contains all the propositions that center believes in that world plus a set that are merely supposed by the center in that world). But that doesn't matter here. Since the modal is indexical to $\odot_{i,j}$, the only value that will matter to the proffered content is that for the indexical anchor and the world of evaluation.

¹⁰ I have chosen the term carefully to accord with its attested meaning in standard dictionaries, like the Merriam-Webster on-line: <http://www.merriam-webster.com/dictionary/>. I can imagine someone arguing that the definition I have given must, itself, be decomposed to reflect the basic properties of what it is to suppose something. But I think that the search for basic semantic units for non-logical terminology is misguided. In the end, non-logical words have their meanings through a combination of associations in use and understood relations to the web of underlying concepts. The lexicon really is a web of words.

b'. Nonetheless, he couldn't believe it: Was he really going to marry Constance?

Again, judgments are subtle, but it seems that Mark's conviction about his obligation is less than complete in (2a), so that he still believes there are ways to avoid doing the unpleasant deed.

However, though it seems to go beyond what is strictly believed, supposition still has many of the properties of belief: One cannot consistently both believe p and suppose $not-p$, or believe $not-p$ but suppose that p :

- (4) #It isn't raining. But I suppose it is.
- (5) #It's raining. But I suppose it isn't.

These have much the flavor of Moore's paradox, a point which bears on Yalcin's (2007) paradox, to be discussed below in §4.2. Note that the case is somewhat different in the imperative, as we see by replacing the second sentences in (4) and (5) with their imperative counterparts: *But suppose it is /it isn't (raining)*. We return to the counterfactuality of imperative *suppose* below.

Accordingly, the semantics above predicts that the following is a consistent assertion:

- (6) George must be the murderer, but I can't believe it!

If in the interpretation of (6) we take the anchor for *must* to be ©*, the distinguished center whose agent is the speaker, then asserting *must p* is weaker than simple assertion of the truth of p : The speaker is only committed to saying that p is true so far as s/he knows, as an inference from the available evidence, which is supposed to be true, i.e. consistent with her beliefs, but not necessarily completely accepted. Thus, the inference of the prejacent p is consistent with a lack of commitment of the speaker to p itself—the speaker needn't yet confidently believe it to be true.

In other words, we have derived a meaning that gives rise to a parallel to Moore's paradox, as in (4) and (5), but without belief in (or, hence, knowledge of) the prejacent *per se*.

Note that since anchoring to a center makes the relevant content *de se* from that center's point of view, that predicts that the anchoring agent should *know* that she knows that the prejacent follows from the supposed evidence (i.e. the agent is such that in all the centered worlds in her belief state, the center is aware of the evidence and of the fact that the prejacent follows from the evidence). So we also predict the inconsistency of:¹¹

- (7) #/! George must be the murderer, but I'm not sure that it follows from the evidence available to me.
- (8) #/! {Cissy thinks that/According to Cissy,} George must be the murderer, but she doesn't realize that the evidence available to her entails that he is.

¹¹ The '!' here marks semantic anomaly. The present theory predicts that that is the reason for the intuitively attested infelicity indicated by '#'.

The fact that epistemic modals are perspectival, plus independently motivated assumptions about context, explains and even predicts a wide range of attested properties of *must*, several of which we will explore in the next section. Here, let me only note that the above semantics does not *entail* that the evidence which entails the prejacent is indirect or in any way inferior in quality to the agent’s beliefs. It merely says that the evidence goes *beyond* what she believes, since $\cap S \subset \text{Dox}(\odot_{i,j})(w)$, i.e. the set of propositions S given by the Modal Base f **properly** entails the anchoring agent’s beliefs at the relevant time. This naturally gives rise to a Quantity implicature.

Most often, if one doesn’t yet accept some proposition p as true, i.e. doesn’t believe that p , yet has explicitly considered whether p (as is necessary in order to suppose it), that would be because the evidence one had for the truth of p wasn’t of sufficiently high quality to foster conviction. For example, the evidence might be merely circumstantial, leaving open other possibilities, even if less likely; or it might involve hearsay, which in turn requires one to have confidence in the source of the report; or it might be based on reasoning to the best explanation, which is also indirect, etc. Any of these might explain why one *merely supposes* p instead of properly believing it. A speaker is committed to believing what she asserts. So asserting that p (based on purported belief) is stronger than merely asserting that it follows on the basis of what one *supposes*. Hence, so long as p is relevant to the Question Under Discussion (so that one *should* proffer it if one truthfully can and is cooperatively committed to resolving the QUD), *must p* conversationally implicates that one is not in a position to assert p , i.e. that one does not believe it. And this is the source of the sense that *must p* is “weaker” than p alone. Since this is a conversational implicature, one would expect that insofar as what it is to be an epistemic modal is to be based on supposition—that these are really *suppositional modals*—we should find the same implicature across languages, as von Stechow & Gillies (2010) claim.

Another way of putting this is to say that epistemic modality is evidential. We clearly owe this insight to von Stechow & Gillies (2010). But unlike von Stechow & Gillies, I assume here, and argue in §4.1, that evidentiality is suppositional. I just suggested that there are a variety of motives to merely suppose instead of believing that p . In some languages, evidentials are specialized to target specific reasons for merely supposing: e.g. in Quechua, we have both reportative evidentials (indicating that p is supposed on the basis of hearsay) and inferential evidentials (indicating that p is supposed on the basis of inference). And we have other flavors of evidentiality across languages. EMAs like *must* and *may* are more general, less specific. But they reflect what it is to be evidential in the general case, and hence we expect similar behavior across languages, including the display of “weakness” relative to straightforward assertion of the prejacent, not only with EMAs, but with evidential particles. So far as I know, this is the case.

Now we turn to a more detailed exploration of the predictions of this semantics for *must*.

4. Application to the puzzles

4.1. The modal base for *must* and *might*

In this section we explore the implications of the type of modal base proposed for *must* and *might* in §3. We consider the frequent claim in the literature that the necessity associated with

must is weak (§4.1.1), briefly consider some ways in which the modal base of *must* differs from those for epistemic *should* and *ought to* (§4.1.2), and briefly look at the role of the QUD in determining the intended modal base of epistemic modals (§4.1.3).

4.1.1 **Weak necessity**

A well-known apparent weakness in epistemic modality is reflected in the pattern illustrated by von Fintel & Gillies (2010) examples (9) and (10):

- (9) [Seeing the pouring rain]
 a. It's raining.
 b. ??It must be raining.
- (10) [Seeing wet rain gear and knowing rain is the only possible cause]
 a. It's raining.
 b. It must be raining.

In such minimal pairs, asserting *must p* is consistently infelicitous when one would be in a position to simply assert *p* instead. Some have argued that this would be unexpected if we take *must* to have the force of simple necessity, since with unrestricted modality $\Box p$ entails *p* (see the useful overview of these matters in Portner 2009, Chapter 4). von Fintel & Gillies (2010) provide a compelling argument that English epistemic modals like *must* and *may* are strong but evidential, with the evidence in question indirect, this feature of their character then explaining the apparent weakness. I refer the reader to their arguments, which I take to be convincing.

More generally, they claim:

We have not found a language whose expression of epistemic necessity fails to carry **an evidential signal of indirect inference**. That is, the paradigm illustrated for English in [(9) and (10)] can be replicated in language after language. This should raise the suspicion that what we are dealing with should not be a stipulated, arbitrary part of the lexical meaning of epistemic necessity modals, and so it shouldn't be a lexically specified presupposition or conventional implicature. Rather, one would suspect and hope that the evidential signal can be derived as a predictable conversational implicature that is non-detachable in Gricean terms. (2010:367)

In their semantics, the indirect evidentiality of *must* is directly stipulated, as a conventional presupposition:

Def'n 4: **Kernels and bases**: K is a kernel for B_K , B_K is determined by the kernel K , only if:

- i. K is a set of propositions (if $P \in K$ then $P \subseteq W$).
- ii. $B_K = \bigcap K$

Def'n 5: **Strong *must* + evidentiality**. Fix a *c*-relevant kernel K :

- i. $[[must \varphi]]^{c,w}$ is defined only if K does not directly settle $[[\varphi]]^c$
- ii. If defined, $[[must \varphi]]^{c,w} = 1$ iff $B_K \subseteq [[\varphi]]^c$

Def'n 5 presumes that K doesn't directly settle φ , but proffers that B_K entails it.

We...see no choice but to stipulate the evidential component of *must* in its lexical semantics, and we have to leave as unsolved the mystery of why this seems to be happening with every epistemic necessity modal that we have come across. (2010:368)

Comparing von Fintel & Gillies' semantics with that offered in §3, we see that though both proposals take the EMAS to be evidential, theirs differs from my own in several respects. The central difference, I would argue, is that the characterization of supposition offered in §3 takes it to be related to belief, so that the modal base of the auxiliaries is doxastic; while von Fintel & Gillies do not appeal to belief, let alone the beliefs of any particular agent. This is at the core of the following important features of my proposal, which differentiate it from theirs:

- (a) Rather than a "cloud of admissible contexts" (von Fintel & Gillies 2008) which would leave the anchor unspecified, I assume that in felicitous use of a context-sensitive expression like *must* the interlocutors have access to a context of utterance which makes available a limited range of discourse centers and makes clear what's at-issue, hence RELEVANT.
- (b) The Character of *must* presupposes that it is anchored to one of the contextually salient discourse centers and thereby to the beliefs of a particular agent whose doxastic state is relevant. As with any anaphoric presupposition, felicitous use requires that this presupposition can be readily resolved in the context of utterance.
- (c) The suppositional modal base requires consistency with the beliefs of the anchor. But also:
- (d) The conditions on the Modal Base require that the evidential ground *S* supporting the prejacent *p* go properly beyond the anchoring agent's beliefs. Hence:
- (e) This semantics conversationally implicates that the evidential ground is consistent with but qualitatively weaker than that of the agent's firm beliefs, as sketched in §3.

This last difference, (e), I take to satisfy the desideratum they note in their quote above, that the evidential signal should be derived as a non-detachable, predictable conversational implicature, of the sort sketched in §3. Like von Fintel & Gillies, I do not weaken *must* through use of an ordering source (Kratzer 1991), or by strengthening the prejacent by making it a test on the context (Veltman 1985); nor do I treat EMAS as speech act operators (e.g. Lyons 1977; cf. Faller 2002 on Quechua evidentials). The modal force is that of simple necessity, and the apparent weakness comes not through the OS but as a consequence of implicated qualitative weakness of the evidential base. The conclusions one reports with it are weaker than beliefs: The evidence isn't strong enough to confidently simply assert *p*. This implication of qualitative weakness, since it is a quantity implicature, should arise in the same way cross-linguistically for all evidentials (and not just EMAS), though evidentials are often specialized to presuppose a particular type of supposed evidence (inferential, hearsay, etc.). One *supposes* that what one has heard, inferred is correct, but there's room for error, so no firm belief.

Glass (2013) argues that von Fintel & Gillies' indirectness requirement is incorrect, and that, instead, epistemic *must* does have weak readings, and (roughly) merely requires that the prejacent is *inferred* from the premises given by the modal base. Her argument against indirectness hinges on examples like the following:

- (11) The answer is divisible by 2 within no remainder, so it must be even.

which, Glass argues is felicitous despite the fact that the speaker has absolutely no doubt of the truth of the prejacent—here, that *s* is divisible by 2. But I do not share her intuitions about this example. To me it's no better than a person standing in the pouring rain saying *it must be raining*. But compare:

(12) For arbitrary *x*: if *x* is even, it must be divisible by 2 with no remainder.

This is minimally different from (11), but much better for me. I suspect that this is due to the arbitrary nature of *x*: In (11), one can take the (presumably specific) answer in question and inspect it, perform a concrete calculation, and thereby ascertain that it is divisible by 2, and therefore (by definition) even. I.e., having calculated its divisibility, its evenness is a matter of definition. But in (12), the number is arbitrary, so we know nothing of it at all. Then we hypothetically assume that it's even, from which (again, by definition) it follows that it's divisible by 2. Notice that the hypothetical assumption in the *if* clause is crucial here, and that such an assumption is not a *belief*—it goes beyond what the addressee already knows/believes about the arbitrary *x*. So this is consistent with the constraints presupposed by *must* in §3.

4.1.2 Differences in admissible modal bases for different epistemics

The Kratzerian approach to modals tends to emphasize the fact that the proffered content of a given modal auxiliary is indeterminate with respect to its modal base and ordering source, which are given by context, contributing only the force of human necessity or possibility. But it is well-known that not all modal auxiliaries can yield all types, or flavors, of modality (Kratzer 1981,1991; Palmer 1990). For example, *must*, *would* and *needn't* all have the force of necessity. But while *must* may have either an epistemic or a deontic interpretation (and not, for example, a circumstantial flavor), *would* typically has a circumstantial flavor, often counterfactual, and in contemporary English *needn't* has only a deontic flavor. *May* and *can* both have the force of possibility, but like *must* and *might*, *may* may be either epistemic or deontic, while *can* has a deontic or dynamic flavor.¹² How might these differences be captured?

Note that these distinguishing features of the relevant modals are very much like lexical subcategorization features, the latter constraining the sort of arguments that a head/functor can take. Above I take them to be captured in the Presupposed content of the modal auxiliaries' Characters. But note the particular constraints on the modal base of *must* or *might* are lexical, not necessarily intended to suffice in characterizing all the epistemic modals. One possibility is that what it is to be epistemic is to be evidential in the sense given here, but that other epistemic modals have additional constraints. Here is an illustration of what I mean, noticed by Stone (1994). Kratzer (1991) had argued that epistemic modals take a stereotypical ordering source, the defeasibility of information about what is typically true accounting for the modals' apparent weakness. Stone considers the following examples (pp.3-4), which argue that this works for *should* but not for *must*:

¹² *can* is sometimes said to have an epistemic interpretation, as well, though I am somewhat skeptical of that claim, since the candidate cases I've considered carefully involve something very like logical possibility, which is presumably circumstantial. But these are sometimes very subtle differences.

- (13) John: Where is the sugar?
Bill: It should be in the cabinet over the fridge.
- (14) Ann: Where is the sugar?
Mary: It must be in the cabinet over the fridge.

Stone's discussion:

Imagine stumbling across either of these sentences as part of a story. Understanding them involves explaining how come the first speaker asks this question, and how come the second arrives at this answer. When the discourse is presented out of context, this interpretation triggers the accommodation of several assumptions. For example, you can infer from John's and Ann's question that they are carrying out some activity for which they need sugar and that they find themselves in circumstances where they expect their addressee to know where the sugar is.

The assumptions accommodated for Bill's reply are quite different from those accommodated for Ann's. Bill's reply seems to invoke a modality of epistemic necessity along the lines suggested [by Kratzer]: Bill answers that, if everything is normal, in view of his knowledge, the sugar is in the cabinet over the fridge. To understand Bill's utterance, the reader need only accommodate some reason for Bill to know where the sugar normally is. You might suppose that John is at Bill's house making cookies (this explains the question) and that Bill normally puts the sugar over the fridge, but sometimes accidentally exchanges it with flour which goes in the cabinet over the stove.

In contrast, to make sense of Mary's answer in [(14)], one must assume that Mary has just seen something or figured something out from which she concludes that the sugar is in the cabinet over the fridge. Perhaps Mary has seen a telltale trail of white particles, or perhaps she has realized that only one cabinet remains in the kitchen which Ann has not ruled out. . .

The implicatures of Bill's response are also quite different from those of Mary's. Bill's utterance carries an implicature that Bill does not know whether the world is normal: perhaps he cannot remember whether this is one of the days he accidentally put the sugar over the stove. Indeed, suppose John subsequently looks over the fridge and finds no sugar. If this happened, could John then reproach Bill for the bad directions? My intuitions say no. Bill could respond that the sugar does indeed normally go in the cabinet over the fridge, that this instance is not normal, and that he warned John that this might be the case.

To make her reply in [(14)], meanwhile, Mary must believe that the sugar is in the cabinet over the fridge. There is no suggestion that the sugar will be there only if things turn out normally. For example, should Ann open the cabinet and find nothing, Mary would be proved wrong. Mary could not then say, "I was correct to infer that the sugar was in there, but this is just an abnormal circumstance."

Stone concludes that "the pattern of accommodation and implicature for [(13)] is precisely what is predicted from the assumption that Bill's use of *should* semantically marks the information he presents as **tentative or defeasible**" [my emphasis].

This seems just right to me. I don't think we can take the difference to argue that the modal base of *must* altogether excludes considerations of how things normally go. Rather, *should* (and, I suspect, *ought to*) seems to presuppose that such considerations are important support for the truth of the prejacent. That being the case, if stereotypicality is crucial support for the prejacent, it would be clearer to use *should* (or *ought to*) rather than *must*. But I won't attempt to develop this argument here; for extended discussion of *should* and *ought* see von Stechow & Iatridou

(2008), where they call these “weak necessity modals” and offer strong cross-linguistic evidence for their character.¹³ The point is that though both *must* and *should* have modal bases that go beyond what the anchoring agent strictly believes, the latter seems to have an additional requirement that what is supposed is defeasible.

This is all part of a much richer story to be told about the subtle differences in meaning between different English modal auxiliaries, differences which I suspect are partly a function of their different etymologies. For example, Ramchand (2014) notes a fascinating difference between *must* and *might*: the former only has epistemic interpretations when its prejacent displays stative aktionsarten, while latter permits dynamic, non-stative aktionsarten, as well. Explaining this difference would take us too far afield from the present discussion,¹⁴ but it illustrates how, though the constraints on admissible modal bases for *must* and *might* proposed above may be necessary, that does not guarantee that they are sufficient even for these two epistemic modals, let alone for all auxiliaries displaying that modal flavor.

4.1.3 **The QUD and the Modal Base**

Moss (2015) provides extensive evidence for the central role of the QUD in the determination of the intended modal base for an epistemic modal. Of particular interest here are a set of examples (inspired by examples discussed in Lycan 2001, Slote 1978) involving epistemic *probably*, in the following scenario:

Jill is standing on the roof of your office building. The local fire department occasionally hangs a net along the roof to protect workers doing construction. The net is strong enough to safely catch anyone who falls off the building. Just a few hours ago, you happened to notice that there was no net along the roof. As a result, you do not believe that Jill is going to jump off the roof. Jill is a thrill-seeker who might jump into a net for fun, but she definitely does not have a death wish. And without a net, anyone who jumped off the roof would surely fall to the ground and die. [Moss 2015:16ff]

Moss argues convincingly that what’s assertable about this scenario, on the basis of the evidence just described, depends on the question being addressed. On the one hand, suppose that the topic of discussion is the circumstances on the roof:¹⁵

- (15) A: Is there a net on the roof?
B: If Jill jumps off the building, she will probably die.

¹³ Here is an actually occurring café restroom sign reported by von Fintel & Iatridou:

(i) After using the bathroom, everybody ought to wash their hands; employees have to.

As they point out, in this example we can substitute *must* for *have to*, *salva veritate*. The clear sense is that *have to* or *must* is stronger than *ought to*.

¹⁴ though I would argue that it can be done without the radically different modal semantics proposed by Ramchand.

¹⁵ I have slightly modified Moss’ examples to make them sound most natural, without, I think, making any changes essential to her argument. My discussion describes the examples in a way that I think reflects her intentions, without parroting her own exposition. So refer to the original before criticizing Moss.

If B's reply is to be understood as RELEVANT (in the sense defined in §2) to A's question, then it must be taken to be indirect, for what Jill does is only RELEVANT insofar as it entails something about whether there's a net. Then if we understand B to mean something like 'given what I know about the circumstances on the roof, its height, and how far a human can fall and still survive, it is likely Jill would die if she jumped'. Jill's character is irRELEVANT here.

But suppose instead that the topic of conversation is Jill herself:

- (16) A: Is Jill suicidal?
B: If Jill jumps off the building, she will probably live.

Again, B's reply to A's question is indirect, RELEVANT only insofar as it contextually entails information about Jill's character, for which the properties of the roof itself are only RELEVANT insofar as they provide evidence about that character. So we might understand B to mean 'given what I know about Jill's character (her joy in life and her intelligent awareness of her surroundings) and the strength of the net that's usually on the roof, etc., if she chooses to jump, it is most likely that she'll live'.

Thus, the requirement of RELEVANCE to the QUD places a strong constraint on the type of evidential modal base we can take the speaker to presuppose for the interpretation of *probably*. The QUD defines what's at-issue, and that, in turn, constrains what we can take the speaker to intend the modal base to be if the utterance is to be RELEVANT to what's at-issue.

Moss then extends the examples about this same kind of scenario to consider a case where one of the target conditionals in question is embedded in the antecedent of an indicative conditional. Suppose three by-standers are talking about the situation on the roof, seeing Jill poised precariously on the edge. C has a set of binoculars, so has better evidence about the situation than A or B:

- (17) A: Is there a net on the roof?
C: Jill will probably live if she jumps.
B: If it's the case that Jill will probably live if she jumps, then there is a net.

B's assertion would be warranted in this case on the assumption that that C's reply is intended to be RELEVANT to A's question and that C's evidence serves as the modal base for epistemic *probably*. But suppose that these interlocutors had just a few minutes before had the conversation in (16), and that all had accepted (16B) as true:

- (16) B: If Jill jumps, she will probably live.

Would A, B, and C be warranted now in concluding from (16B) and (17B) that there is a net, as an instance of Modus Ponens? Consulting our intuitions, we find that they would not. And the reason is clear: RELEVANCE to the QUD led to the assumption of different modal bases to restrict the domain of *probably* in the two utterances. Hence, though (inverted) (16B) and the antecedent of (17B) are identical on the surface, they denote distinct propositions. Thus,

assuming that non-synonymy should somehow be reflected in the logical forms of the premises in an argument, Modus Ponens does not apply here.

Moss argues for a very different type of semantics for many instances of epistemic modals, basing this proposal on evidence from gradable adjectives like *probably* and their interaction with each other and with epistemic modal auxiliaries like *must* and *might* in cases where they are nested in multiple embeddings. I won't address her proposal directly here as, again, it goes well beyond the central issues addressed in my own. But Moss does assume that some epistemic modals sometimes have their standard semantics. The point here, again, is that both on their standard semantics and presumably in the more complex and graded cases that form the core of her argument, epistemic modals are sensitive to contextual factors in several ways, the QUD being one of them. Below we'll see cases where the presuppositions of *must* and *might* also bear on cases discussed by Moss. Thus, these types of context-sensitivity should be reflected in whatever semantics one ultimately adopts.

In the examples considered so far in this section, the tacit assumption has been that the evidence in the modal base of the EMA is that of the speaker. But this is not always the case. So now let us turn to another central feature of the semantics and pragmatics of EMAs: the way in which they are variably anchored, as a feature of the contextually available discourse centers. Then in §4.3 I will show how the fact that the evidentiality of EMAs is indexically anchored (a) and requires consistency with the anchoring agent's beliefs (b) helps to illuminate a related puzzle about epistemic modals and belief, due to Yalcin (2007).

4.2. Variable anchoring

There has been a good deal of discussion and debate about whose epistemic state might be appealed to in the semantics of epistemic modals, e.g. in Hacking (1967), DeRose (1991), Egan et al. (2005), MacFarlane (2006), Stephenson (2007), and von Fintel & Gillies (2007a). As Yalcin (2007) says, "It is a striking fact that these questions do not have obvious answers." Everyone notes that there is at least a default tendency to understand the speaker to be the relevant agent whose epistemic state is at-issue. DeRose (1991) proposes a *speaker inclusion constraint* to capture this in the strongest way, *requiring* the speaker to be at least part of the group with the relevant state:

speaker inclusion constraint: the relevant community [must] include the speaker.

Hence "whenever S truly utters *a might be F*, S does not know that *a* is not *F*." (p.5)

But as discussed by von Fintel & Gillies (2007a), this requirement is much too strong. Here is a sample of the types of anchors attested for epistemic modal auxiliaries:

(a) speaker at Utterance Time (©*):

(18) John might be the thief.

(19) This suggests that Angela must be in Austin right now.

Especially out of context, with no explicit suggestion to shift to a different point of view, this is the only reasonable way to anchor an EMA: the speaker's point of view is always the default, for all perspectival expressions. And the following is also consistent with DeRose's proposal:

(b) Group containing speaker at some actual time \neq Utterance Time:

(20) Given what we knew at the time, John might have been the thief. [von Fintel & Gillies]

Here the adverbial explicitly shifts the doxastic point of view to that of the speaker and others in the anaphorically retrieved denotation of *we* at some time *prior* to the utterance time.

But we also see cases where the grounds appeal to appear to be solely those of the addressee:

(c) Addressee:

(21) Where might you have put the keys? [von Fintel & Gillies]

As with evidential particles in many languages, this amounts to Interrogative Flip (Faller 2002)—wherein the anchoring perspective naturally becomes that of the addressee when seeking information they might offer.

(d) Arbitrary group containing addressee at some hypothetical time:

(22) [Military trainer:] Before you walk into an area where there are lots of high trees, if there might be snipers hiding in the branches use your flamethrowers to clear away the foliage. [von Fintel & Iatridou 2003]

In (22) note the imperative mood and the dependent plural *your flamethrowers*. This constitutes generic instructions issued to the addressees. Then the *if*-clause constitutes a precondition for carrying out the instructions—they are for circumstances in which the relevant information leaves open the possibility that there are snipers. But then, of course, the only reasonable anchor for *might* would be the trainee who has to decide whether the correct conditions obtain.

(e) Explicit third person orientation:

(23) As far as Bill knows, John might be the thief. [von Fintel & Gillies]
(24) From John's point of view, it must be raining.
(25) This suggested to George that the polar ice cap might be melting.
(26) John thinks it must be raining.

(23), (24) and (25) illustrate several ways to explicitly shift the intended point of view. Note that the center's time may be the present (as is the only possibility with temporally present *must* in (24)), or may be some past time (as in (25), the past suggestion-time). With attitude predicates, like *think* in (26), shifting to the agent is the default (Stephenson 2007). Note, however, that this may not be a necessary shift. For example, in (27) and (28) the anchor probably includes the speaker (with or without the addressee), though perhaps a group that

includes both the speaker and John, but the latter hasn't yet drawn the relevant conclusion from the information to which he's privy:

- (27) John won't acknowledge/hasn't yet conceded/hasn't realized that it must be raining.
- (28) Has John realized that it must be raining?

Jefferson Barlew (p.c.) offers a case where the anchor would not include the agent of the attitude, here the speaker:

- (29) Bill claims to have discovered evidence allowing him to deduce who killed Frank. If that's true, I hope it must be George.

The only reasonable, RELEVANT interpretation to give *must* in (29) is as anchored to Bill and his new evidence. If that seems a bit forced, it just illustrates how strong the tendency is to anchor the EMA to the agent of the attitude *hope*, here the speaker. Dowell (2011) notes that it's certainly possible to shift to an agent other than the agent of the embedding attitude when we use explicit shifting adverbials, as in this variation on one of her examples:

- (30) Leiter believes that, for all Blofield and No.2 know, Bond might be in Zurich.

and she gives other convincing cases without explicit adverbial shifting (see her discussion of extensions of the Mastermind example p.23ff).

(f) Third person in extended modal subordination:

- (31) Suppose you were John. Where would you go now to find Clarissa? You might find her with Sidney in New Orleans, or maybe you would find her with her aunt in Chicago.

The supposition proposed to the addressee is to shift his point of view to that of John. The sequel contains a series of modals: *would*, *might*, *maybe*, each of them understood relative to the point of view of John at that point. The disjunction suggests alternative possible answers to the QUD introduced by the preceding interrogative, *might* and *maybe* enumerating the possibilities epistemically accessible to him.

(g) Third person in Free Indirect Discourse (FID):

- (32) John pondered his situation. Where was Clarissa now? She might be in New Orleans with Sidney. But she might be in Chicago.

Here, as typical in FID and contrasting with (31), there is no explicit suggestion to the addressee to shift point of view from that of the speaker to that of the reported agent, John (Eckardt 2014). Assume that the novel takes place in the 19th century, but we are reading it in the 21st. Then the use of *now* in querying Clarissa's location (e.g., if) suggests the shift.

(h) Multiple bodies of evidence, one more "objective" than the other:

- (33) Given the results of the DNA tests, John might be the thief. But if we take the eyewitness seriously, John can't have been the thief. [von Fintel & Gillies]

The speaker in (33) explicitly suggests first one body of evidence, the DNA tests, and then another, the eyewitness' account, which make contradictory predictions about whether John might have been the thief. This is a case where the speaker (and probably the addressee, as well) knows the DNA evidence, and that it can be interpreted to so that John might have done it. The *if*-clause in the last sentence suggests that the speaker and addressee haven't yet decided whether to take the eyewitness' evidence seriously; then *can't* in the main clause reflects the consequences of the doxastic state they *would be in* if they were to suppose that that evidence is correct.¹⁶ Another possibility is that *can't* is anchored to the eyewitness himself, as would be explicit in *But according to the eyewitness*. It may not make a truth conditional difference which anchor we chose for *can't*, since the same body of evidence would be brought to bear on the interpretation in both cases. But the evidential semantics of EMA *can't* would argue for the speaker(+addressee) anchor: If the eyewitness saw the theft, that person would have the best available evidence and the speaker could simply report, *But according to the eyewitness, John was the thief*. So here the EMA is used because from the interlocutors' point of view the evidence is hearsay. This example is especially nice because it shows how context not only suggests the identity of the intended discourse center anchor, but in some cases may suggest the particular body of evidence which is supposed by that agent, going beyond his firm beliefs.

von Fintel & Gillies (2007a) also suggest that objective epistemic modals might be anchored to non-human knowledge sources (logs, charts, etc.), and that in such cases this leads to a more "objective" interpretation where the speaker's opinion doesn't count:

(34) The hulk might be in these waters. [von Fintel & Gillies 2007a, after Hacking 1967]

But whatever the knowledge sources for (34), if the speaker knows of them, and the propositions they make true, then *s/he's* privy to that evidence. If there's an implication that the body of evidence is "objective" or "consensual" or generic, then the modal statement may be taken to have more force—after all, evidentiality is about the speaker's judgment of the quality of her evidence for certain beliefs, and we are more likely to believe that for which we have better evidence. Such examples seem particularly prone to the 'we' interpretation we saw earlier, which arises in a group where the Common Ground contains such "objective" evidence; and I suspect the common consensus about the value of the information strengthens the sense of objectivity, as well. But (34) is still evidential, and the non-human provenance of the evidence doesn't mean that the anchor itself—the epistemic agent who has that information—is non-agentive.

Another type of content in which modal auxiliaries are anchored in the non-default situation are the type we see in von Fintel & Gillies' (2007a,2010) Mastermind examples, which we'll

¹⁶Note (Palmer 1990) that *can't* has an epistemic interpretation, while *can* typically does not. And that epistemic *can't* always has scope order $\neg\Diamond$, whereas epistemic *mightn't* (as in *Mightn't George be the thief?*) has only $\Diamond\neg$. Thus, epistemic *can* arguably is a Negative Polarity Item (probably of a particularly restrictive sort, requiring negation), in its epistemic interpretation (but not in the deontic or dynamic) requiring a downward entailing context for felicity.

consider further below, in §4.5, where we'll also discuss some other examples that hinge not only on what the anchoring center is, but on what's at-issue.

We'll consider in more detail below other examples of how the anchoring agent is established in a particular context and of its role in interpreting the corresponding EMAs. Here let me reiterate the central hypothesis about discourse centers: For an individual to be a discourse center, it does not suffice that that individual be an agent who is familiar and salient to the interlocutors and is capable of having a doxastic point of view. Instead, that individual's doxastic state itself must be relevant. Hence, the restriction of doxastic anchors to the set of discourse centers is intended to constrain the range of possible EMA anchors in a given context of utterance.

The pattern of occurrence observed above is that typical of anchors for indexical expressions generally (Roberts 2015), wherein the presupposed center who anchors the indexical's interpretation varies systematically with the context. With EMAs, not only does the anchor vary, but it then gives clues about the body of evidence which is intended to restrict the modal's domain. Summarizing, we see that for EMA domain restriction:

- in root declaratives, evidence to which the speaker or inclusive 'we' is privy is the default
- in root interrogatives, evidence to which the addressee is privy is the default, i.e. interrogative flip (21), though as we'll see in some cases to be discussed in §4.5, that is not necessary
- in imperative conditional protases, as in (22), evidence to which the addressee is hypothetically privy is the default, as this information is useful for determining the applicability of conditional advice or directions
- in attitude complements, evidence to which the attitude agent is privy is the default
- with explicit adverbials like those in (23),(24) and (25), we see a shift similar to that in attitude complements.
- in modal subordination the evidentiality is anchored to the relevant agent in the epistemic subordinating context
- in FID the default is the evidence of the doxastic agent whose perspective is adopted by the author.

In almost all of these cases, the default can be pragmatically over-ridden. The resolution of the anchoring presupposition of an EMA is anaphoric, hence subject to the usual constraints on anaphora resolution: familiarity and salience of the intended discourse referent antecedent, pragmatic plausibility, and coherence of the resulting resolution. A useful contrast is between the anchoring of EMAs in attitude complements and those governed by explicit shifting adverbials. With the adverbials, the shift appears to be obligatory, conventional, whereas it is only the pragmatically governed default in the attitude complements. In the latter it is just that the agent of the attitude is the most salient center in that context, so that anaphora resolution naturally points to that center, all other things being equal.

Variable anchoring is closely related to another important issue in the recent literature on epistemic modality: the argument from faultless disagreement for modal Relativism. We will take that up in §4.3.1 and again in §4.5.1 below.

4.3 Epistemic modals and attitudes

In the recent literature, proofs involving semantic interactions between EMAS and attitudes like knowledge and belief have been used to argue for Relativism—the post-truth conditional assessment of content relative to a Judge—and for a kind of Expressivism about EMAS—a failure of the clauses in which they occur to have their usual truth conditional content. Here I will consider two prominent accounts along these lines, due to Egan et al. (2007) and Yalcin (2007), and argue that on the present account we need neither Relativism nor Expressivism to predict the kinds of interpretations attested.

4.3.1 Granger’s puzzling proof and the *de se* Character of *must* and *might*

Determining the intended anchor for a given modal auxiliary not only makes a difference to truth conditions, but, as we might expect, also to the patterns of inference associated with the use of epistemic modal auxiliaries. This is especially evident in examples where the anchor shifts across different uses of the same modal in a single argument, given the typically privileged status of an anchoring agent’s evidence. Moreover, the pragmatic basis of the determination of the anchor for an epistemic modal auxiliary differs from the way that the anchor for an attitude predicate like *know* is determined compositionally, though both operators are *de se* with respect to the anchoring agent. So when all these factors interact, the result can lead to some complexity in analyzing the resulting argument.

We see all this in the following example from Egan et al. (2005), which is the basis of their central argument for relativism about epistemic modals:

- (35) [Context: Professor Granger is in the South Pacific and knows it. She heard her friend Myles, in Boston, on the radio speculating about where the missing Granger might be. Myles knew that she’d originally planned to go to Prague, but not where she actually went. She puzzles about the following statements, each apparently true:]
- (1) When he says, “She might be in Prague” Myles says that I might be in Prague.
 - (2) When he says, “She might be in Prague” Myles speaks truly iff neither he nor any of his mates know that I’m not in Prague.
 - (3) Neither Myles nor any of his mates know that I’m not in Prague.
 - (4) If Myles speaks truly when he says that I might be in Prague, then I might be in Prague.
 - (5) I know I’m not in Prague.
 - (6) It’s not the case that I know I’m not in Prague if I might be in Prague.

Supposedly, in (35) steps (1) – (5) lead to a contradiction with (6), but that seems counterintuitive. The poor little professor is having a hard time (per Egan et al.) figuring it out. Let’s help her out. Listen, honey:

The puzzle comes from looking at surface form only, giving rise to the following apparently sound argument (taking *I* in (35) to consistently refer to Granger):

- Premises: the evidently true (1) – (6)
- (2)+(3) |= Miles speaks truly when he says “She might be in Prague” = (a)
 - (a)+(1) |= Miles speaks truly when he says Granger might be in Prague = (b)
 - (b)+(4) |= Granger might be in Prague = (c)
 - (6)+(c) |= It’s not the case that Granger knows Granger is not in Prague. = (d)
 - (5)+(d) |= Granger knows Granger is not in Prague and it’s not the case that Granger knows Granger is not in Prague.

Contradiction.

But as we have seen in earlier sections, epistemic modal auxiliaries are indexical, anchored to an available doxastic agent. This means that in assessing truth conditions and propositions expressed, we have to take into account not just the surface form of the premises in the argument, but the way the modals are anchored. Ignoring (6) for the moment, it is clear from the story that the following anchors are understood, where subscript *M* indicates that Miles is the doxastic anchor, *G* indicates that it is Granger:

- (1) When he says, “She might_M be in Prague” Myles says that I_G might_M be in Prague.
- (2) When he says, “She might_M be in Prague” Myles speaks truly iff neither he nor any of his mates know that I_G’m not in Prague.
- (3) Neither Myles nor any of his mates know that I_G’m not in Prague.
- (4) If Myles speaks truly when he says that I_G might_M be in Prague, then I_G might_M be in Prague.
- (5) I_G know I_G’m not in Prague.
- (6) It’s not the case that I_G know I_G’m not in Prague if I_G might be in Prague.

In (1) and (2) as uttered in the story, *might* is clearly anchored to Myles, because Granger is reporting what Miles said when he was talking about the missing Granger and answering questions about what *he* knew about where she might be. The only reasonable way to understand *might* in (4) is with the anchoring indicated. In the protasis *must* has to be anchored to Miles because, as we saw in the abbreviated logical form of the argument above, this is the only way for (4) to combine with (b) to allow us to deduce (c): to license Modus Ponens, it is crucial to have the same anchoring in (b) as in the *if*-clause of (4). And the plausibility of (4) itself depends on it reflecting the natural assumption that if one speaks truly in claiming that *p*, then *p* is true. Hence, the proposition in the consequent must have the same logical form as that in the complement of *says* in the antecedent. And thus, in the logical form of (c), so deduced, *might* is anchored to Miles, as it was in (b) and (4). (5) involves no epistemic modal, but it has the verb *know*, itself epistemic. Given its lexical semantics, *know* has to take its subject’s denotation, Granger as the anchoring agent for the reported belief state: It’s true in all Granger’s belief-worlds that she’s not in Prague.

Then what does (6) mean? For the puzzle to arise, Egan et al. again need this to be a fairly obvious truth, so they need a logical form like the following:

- (6’) if I_G might_A be in Prague [then] it’s not the case that I_G know I_G’m not in Prague the protasis:

I_G might_A be in Prague:

◇_AP(G)

the main clause:

scope of negation:

$I_G \text{ know}_A I_G \text{ 'm not in Prague}$ $\Box_A \neg [P(G)]$

so:

it's not the case that $I_G \text{ know}_A I_G \text{ 'm not in Prague}$ $\neg(\Box_A \neg [P(G)]) = \Diamond_A P(G)$

This logical form requires that the same agent that anchors the doxastic state appealed to in the lexical content of *know* also anchors *might*, so that the content of the two clauses is essentially equivalent. But in the story, the speaker Granger is the denotation of the subject of *know*. Hence, to get the plausible logical form, we have to take Granger to anchor *must* in the protasis: $A = \text{Granger}$. Then (6') means 'if for all Granger knows Granger might_G be in Prague, then it's not the case that Granger knows she's not in Prague'.

But plugging this back into (35) will not yield the contradiction. In the schematic derivation, that contradiction was entailed by (6) plus (c) under Modus Ponens. MP is only applicable if (c) has the same logical form as the protasis of (6). But we have seen that in (c) *must* is anchored by Miles, while in (6), it is anchored by Granger.

For MP to apply, we would need the following logical form for (6):

(6'') if I might_M be in Prague [then] it's not the case that I know_G I'm not in Prague
 $I_G \text{ might}_M \text{ be in Prague:}$ $\Diamond_M P(G)$
it's not the case that $I_G \text{ know}_G I \text{ 'm not in Prague}$ $\neg(\Box_G \neg [P(G)]) = \Diamond_G P(G)$
 $\neq \Diamond_M P(G)$
'if for all Miles knows Granger might be in Prague, then it's not the case that Granger knows that Granger is not in Prague'

But (6'') simply isn't just true. Just because Miles is in the dark, that doesn't mean that Granger doesn't know where she is. And the force of the argument depended on the plausibility of each of the premises, including (6).

So, on the only reasonable interpretation of (6), (6'), there's no contradiction with (5).

The crucial premise here, then, was (6). And it was important for its plausibility that the epistemic agent of the modality implicit in *know* was the same as the anchor of *must*. But though both *know* and *must* are doxastic and *de se*, they differ in how they may be anchored: *must* is variably anchored, the intended anchoring contextually resolved, while the agent of *know* is given lexically and *must* always be its subject.

This gives rise to another set of possibilities for the interaction between attitude predicates like *know* and EMAS, one which also highlights the *de se* nature of the anchoring. When attitudes and EMAS are iterated, there are a variety of possible readings for the embedded EMAS, just as the present account would predict, hinging on how their anchoring relates to that of the embedding predicates.

Consider the following extension of the scenario entertained in (35):¹⁷

Granger has passed out from drinking too much kava kava. When she awakes, she doesn't remember who she is and her passport is missing. Taken to a hospital, she's been listening to the news about the missing Professor Granger, who according to this fellow Miles on the news was on her way to Prague. A nurse comes in and asks her the question in (36):

(36) Q: What have you heard about Granger's whereabouts?

Then in that doxastic state, Granger might answer (37), but it would be odd—and in fact false, for her to answer (38):

(37) $I_{\odot^*=G}$ know that Granger $must_{\odot^*}$ be in Prague.
 $know(\langle \odot^*, w^* \rangle) \subseteq \{ \langle \odot, w \rangle \mid [\bigwedge \langle \odot, w \rangle] \subseteq \{ \langle \odot', w' \rangle \mid \underline{\text{Granger is in Prague in } w'} \} \}$
 where $[\bigwedge \langle \odot, w \rangle] \subset \text{Dox}(\odot)(w)$

since $\text{anchor}(I) = \text{anchor}(must)$, $\odot^* = \odot = \odot'$

[it's not really true that the anchors are in any intuitive sense equivalent; but because there *is* a transitive epistemic self-location relation between them, by virtue of the transitivity of epistemic accessibility relations, Granger self-locates both as the knower and as the person in Prague.]

(38) $I_{\odot^*=G}$ know that I_{\odot^*} $must_{\odot^*}$ be in Prague.
 $know(\langle \odot^*, w^* \rangle) \subseteq \{ \langle \odot, w \rangle \mid [\bigwedge \langle \odot, w \rangle] \subseteq \{ \langle \odot', w' \rangle \mid \underline{\odot' \text{ is in Prague in } w'} \} \}$
 where $[\bigwedge \langle \odot, w \rangle] \subset \text{Dox}(\odot)(w)$

In the logical forms for these examples, I have highlighted in yellow those elements that are self-identified counterparts in their respective doxastic states. The base center is, of course, the speaker, \odot^* (who doesn't think of herself as Granger). In the underlined contents characterizing the doxastically accessible centered worlds, we see that in (37) the arbitrary center, the speaker's counterpart, is not self-identified with Granger, whereas in (38) she is. In both, the speaker *is* Granger, so that it happens that *I* is coreferential with *Granger*.

So in (37), the embedded *Granger* happens to be coreferential with *I*, since the speaker is Granger. But only the anchoring of the EMA *must* is *de se*: its interpretation based on what the speaker takes her own evidence (about Granger) to be. But in (38), the *de se* interpretation of *I* guarantees that the speaker takes herself to refer to herself, incorrectly self-locating in Prague according to her own evidence. This is false because she doesn't know that she's Granger.

But now extend the story yet further: The amnesiac Granger hears on the news a report by a purported eye-witness who claims that she saw Granger arrive at the airport in Prague, but that just as Granger was stepping off the plane, she was stunned by an assailant, passed out and was then kidnapped. Then she might report this to the nurse as follows:

¹⁷ Obviously inspired by Morgan's (1970) Ernie Banks story.

- (39) $I_{\odot^*=G}$ know that from her $_G$ point of view, Granger must $_G$ be in Prague.
 $know(\langle \odot^*, w^* \rangle) \subseteq \{ \langle \odot, w \rangle \mid [\cap f(\langle \odot, w \rangle)] \subseteq \{ \langle \odot', w' \rangle \mid \text{Granger is in Prague in } w' \} \}$
 where $[\cap f(\langle \odot, w \rangle)] \subset \text{Dox}(\odot)(w)$

When we have a case like (37), where the embedded subject is non-*de se* despite being coreferential with the agent of matrix *know*, we also get the possibility that embedded *must* can be (at least explicitly) anchored to *that embedded subject* and hence itself non-*de se* wrt the anchoring subject of matrix *know*. This is what we have in (39). Note that the truth conditions of (37), (38) and (39) are all distinct. (39) shows that when embedded under an attitude, *must* can be anchored to someone other than the agent of the attitude, supporting the claim that its anchoring is presuppositional and pragmatic, unlike the anchoring of the modal in *know*, which is lexically given. And it supports the claim that the anchoring of *must* and other EMAs is *de se*, like indexical *I*. In this case, the *de se* subject agent is the speaker, who is Granger and the anchor of *must* also happens to be Granger, but because the latter doesn't *know* she's Granger (doesn't self-locate as the denotation of *she* and the anchor of *must*), *must* itself isn't anchored to the actual belief state of the speaker, but to the evidentially supposed state of Granger.

4.3.2 Yalcin's (2007) puzzle

Yalcin (2007) observes that pairs of examples like the following appear to pose serious problems for a semantics of epistemic modality that assumes a *de se* perspective like that of Stephenson (2007):

- (40) a. Suppose [it is raining but you don't believe it is].
 b. #Suppose [it is raining but it might not be].

He argues that such examples cannot be addressed within the standard relational semantics of the type developed by Kratzer, and instead offers an alternative, "domain semantics" account based on acceptance. But I will argue that on reasonable assumptions about the meanings of the expressions involved, these examples are not problematic for the account proposed here, and in fact are exactly what it would lead us to predict.

The key to explaining this puzzle lies in explicating the meaning of *suppose*, which involves implicit iterated attitudes. Iterated attitudes are a bear,¹⁸ especially when, as here, some of them are *de se* and counterfactual. Nonetheless, we can tease apart the relevant factors involved in giving rise to this puzzle, and thereby account for it without resorting to the kind of non-standard modal semantics that Yalcin adopts.

Here is the claim I will motivate: To realize an imperative direction to suppose that *p* is to entertain a counterfactual doxastic state. In order to realize the directions in (40b), the addressee would have to construct and entertain a counterfactual revision of her doxastic state that is inconsistent in just the way that Moore's (1993) original examples are epistemically inconsistent:

¹⁸ an American idiom meaning roughly 'challenging to wrestle with', but implicating something about blood and potential dismemberment. I am Davy Crockett's great-great-great niece, so I am licensed to use the expression.

She would have to entertain the proposition that *It's raining and it might not be*. Therefore, we get an embedded counterpart of Moore's paradox: Since we cannot rationally entertain an inconsistent state, these directions are anomalous, and cannot be realized. This is very different from (40a), which merely requires that in the revised doxastic state the addressee's counterfactual counterparts are confused.

Now to the details:

These examples involve imperatives. I assume without argument that imperatives have realization conditions, rather than truth conditions—circumstances in which they would count as being realized because the corresponding declarative with the addressee as subject would be true. Further, the realization conditions of imperatives require that the time of realization be at or after the time of issuance of the directive involving the imperative. (See Roberts 2015b for details, but these assumptions are largely in keeping with the accounts of Portner 2007, Kaufmann 2010.) So roughly:

An **imperative** S! with LF [_S ! VP] addressed to *a* at time *t** in context *D* is realized in world *w* just in case: $\exists t \leq t^*: a \in |VP|^D(w)(t)$

The realization of the imperative can only take place at some present or future time *t* at which the imperative property holds of the addressee in the world of evaluation.

Now notice that, at least in the imperative mood the semantics of *suppose* are consistent with supposing counterfactual propositions: Someone can felicitously issue the directive (40a) when the Common Ground entails that it's raining. This often involves subjunctive mood, as in *Suppose you hadn't eaten that piece of strudel! You'd enjoy your dinner more now*. I have not found counterfactual subjunctive complements of *suppose* in the declarative or interrogative moods, though it is possible I have overlooked them.¹⁹ Here I will focus on the counterfactual use in the imperative mood.

Imperative *suppose* denotes a counterfactual attitude toward the complement proposition *p*. *p* itself may or may not be true, but (as in other moods, as well) it shouldn't be a proposition that the actual agent believes. Then the attitude involves entertaining this counterfactual proposition. What is it to entertain a proposition, as opposed to believing it, doubting it, etc.? I leave the notion undefined,²⁰ assuming that that is a question for philosophy of mind and psychology. But semantic intuition (and the fact that *suppose...* is so often followed by *then...*) suggests that it involves putting oneself in the position of counterfactually taking the world to be as characterized by that proposition in order to consider what follows. Crucially, to do this one doesn't entertain the proposition alone, in isolation. Instead, one is interested in what it implies about the world *against the background of one's other beliefs about the way the world is*. Of course, since this is a counterfactual attitude, the entertained proposition may be inconsistent with some of one's other beliefs. And one cannot reasonably simultaneously entertain inconsistent propositions, for *ex falso quodlibet*: from inconsistency nothing interesting follows, because everything follows. Thus, crucially, entertaining a proposition (or a set of propositions,

¹⁹ References to relevant work appreciated.

²⁰ Merriam-Webster (<http://www.merriam-webster.com/dictionary/entertain>): 3a: "to keep, hold, or maintain in the mind".

via their intersection) presupposes that it is a non-empty set of centered worlds. Accordingly, to entertain p against the backdrop of one's belief state, first p itself must be consistent, and then one must revise those beliefs to yield a state that's consistent with p but in other respects as much as possible like the way things actually are so far as one knows.²¹

Capturing this similarity requirement, as usual in counterfactual modality, requires a similarity metric over possible worlds, in the present case possible centered worlds. Following Heim (1992), I adopt a very simple notion of similarity, intended mainly as a place-holder for whatever notion is most appropriate in the general case. One is invited to substitute one's favorite similarity metric. SIM takes a proposition p (a set of centered worlds) and a centered world $\langle \odot, w \rangle$ and yields the most $\langle \odot, w \rangle$ -like of the p worlds:

$$\mathbf{SIM}(p)(\langle \odot, w \rangle) =_{\text{def}} \{ \langle \odot', w' \rangle \mid \langle \odot', w' \rangle \in p \ \& \ \forall \langle \odot'', w'' \rangle \in p: \langle \odot', w' \rangle \text{ resembles } \langle \odot, w \rangle \text{ at least as much as } \langle \odot'', w'' \rangle \text{ does} \}$$

In terms of SIM, we define a proposition which is based on the doxastic state of an agent, revised only as much as required to make p true:

$$\mathbf{REVDOX} =_{\text{def}} \lambda p \lambda \langle \langle x, t \rangle, w \rangle \lambda \langle \langle y, t' \rangle, w' \rangle. \langle \langle y, t' \rangle, w' \rangle \in (p) \ \& \ \underline{\exists \langle \odot'', w'' \rangle \in \text{DOX}(\langle \langle x, t \rangle, w \rangle): \langle \langle y, t' \rangle, w' \rangle \in \text{SIM}(p)(\langle \odot'', w'' \rangle)}$$

RevDOX takes a centered proposition p and a centered world $\langle \langle x, t \rangle, w \rangle$ to yield the set of centered worlds $\langle \langle y, t' \rangle, w' \rangle$ in which p is true that are most like $\langle x, t \rangle$'s belief worlds in w .²² As its name is intended to suggest, this is itself the revision of a doxastic state (the perspective of the center $\langle x, t \rangle$ in w) to counterfactually entail p . Then:

Character of imperative *suppose*:

Presupposed:

the denotation of the subject is a doxastic agent a
in a world w and time t of evaluation, the denotation of the complement p , is s.t.
 $\text{Dox}(\langle \langle a, t \rangle, w \rangle) \not\subseteq p$

Proffered:²³

$\lambda p_{\langle s, t \rangle} \lambda x_e \lambda t \lambda w \lambda \langle \langle y, t' \rangle, w' \rangle. \text{ENTERTAIN}[\langle \langle y, t' \rangle, w' \rangle, \text{REVDOX}(p)(\langle \langle x, t \rangle, w \rangle)]$

Utterance of *suppose* presupposes that its subject is a doxastic agent and that the proposition denoted by its complement doesn't follow from the agent's beliefs in the world and time of

²¹ This implicitly presupposes that those beliefs themselves are consistent. I would maintain that our use of *believe* itself presupposes that convenient fiction, as evidenced by what we take to follow from the truth of a belief report. Again, from inconsistent beliefs one can draw no interesting conclusions. The problem of inconsistent beliefs is one for epistemology, not semantics.

²² Whether there always is a unique such set is, again, a question that also goes beyond current considerations.

²³ It might be useful as an exercise to compare this proffered content with the following, which omits ENTERTAIN:

$$\lambda p_{\langle s, t \rangle} \lambda x_e \lambda t \lambda w \lambda \langle \langle y, t' \rangle, w' \rangle. \langle \langle y, t' \rangle, w' \rangle \in \text{REVDOX}(p)(\langle \langle x, t \rangle, w \rangle)$$

Given its arguments, this function would yield the proposition that is true in a centered world $\langle \langle y, t' \rangle, w' \rangle$ in case that world is the way that REVDOX says it is. But this is the doxastically counterfactual proposition that the addressee is told to entertain, not what would be true were the imperative realized.

evaluation. One is not told to suppose something one already believes. The proffered content is a function from the centered proposition p denoted by the complement and the denotation of the subject agent, plus a world and time of realization (these last three constituting a centered world) to a centered proposition: As usual with attitudes in centered world semantics, the centers in the derived worlds are the counterparts of the base center, the latter the addressee at realization time in the actual world. Then in those derived worlds, the agent counterfactually entertains the proposition that is the result of revising the agent's actual belief state at realization time to make p true, $\text{REVDox}(p)(\langle\langle x, t \rangle, w \rangle)$.

There is another, pragmatic consequence of proffering *suppose*: As usual with attitude predicates, considering an attitude leads to the introduction to the set of discourse centers \mathbb{C}_D of the agent of that attitude at the time and world in which it holds. This addition is local, normally (in the absence of an extended supposition context, modal subordination or FID) persisting only under the scope of the attitude predicate. But note that in the semantics we have given for *suppose* there are two attitudes: *suppose* itself and the counterfactual ENTERTAINMENT state: the state in which the supposer has counterfactually revised her belief state to include the proposition denoted by the complement. Accordingly, under the scope of *suppose* two new discourse centers and their corresponding doxastic perspectives are introduced:

- For center $\mathbb{C}^{\textcircled{a}} = \langle a, t^* \rangle$, the addressee at utterance time, we introduce the center shifted forward to the realization time t : $\langle a, t \rangle$. Thus we have $\text{DOX}(\langle\langle a, t \rangle, w^* \rangle)$: the doxastic perspective of the actual addressee a at realization time t in the actual world w^* . Since the second argument of REVDox in the proffered content of *suppose* is $\langle\langle a, t \rangle, w \rangle$, w the world of evaluation, and in the imperative $w = w^*$, the actual belief state of the addressee at realization time will be used in the similarity metric that's part of REVDox , constraining the resulting entertained counterfactual state.
- For center $\mathbb{C}^{\text{ENTERTAIN}} = \langle y, t' \rangle$ in the proffered content of the logical form above, the counterparts of the addressee and realization time, we have the counterfactually entertained doxastic perspective of the entertainer in the counterfactual entertainment world w' , i.e. the value of $\text{REVDox}(p)(\langle\langle x, t \rangle, w \rangle)$. I.e., y 's counterfactual doxastic state at t' is one in which p is true but which is otherwise as similar as possible to x 's beliefs about the actual world at realization time.

Hence, in realizing imperative *suppose*, the addressee will have counterparts $\mathbb{C}^{\text{ENTERTAIN}}$ in the counterfactually revised doxastic state given by $\text{REVDox}(p)$.

Now recall the semantics of *might*, repeated here for convenience:

The Character of English epistemic *might*:

Given an utterance $\text{might}_{i,j} p$ in context D :

Presupposed content:

Indexical anchor: In D there is a discourse center $\mathbb{C}_{i,j} = \langle d_i, t_j \rangle \in \mathbb{C}_D$.

Evidential Modal Base: There is a function f mapping centered worlds to sets of centered propositions (each a set of centered worlds) s.t. $f(\langle\mathbb{C}_{i,j}, w \rangle) = S$, where:

- | | |
|--|---|
| $\cap S \neq \emptyset$ & | 'S is consistent' |
| $\forall p \in S$: $\text{suppose}(p, \mathbb{C}_{i,j}, w)$ & | 'all p in S are supposed by d_i at t_j in w ' |
| $\cap S \subset \text{Dox}(\mathbb{C}_{i,j})(w)$ | 'S properly entails d_i 's beliefs at t_j ' |

Proffered content:

$$\lambda p_{\langle s, t \rangle}. \lambda \langle \odot, w \rangle. p \cap [\bigcap f(\langle \langle d_i, t_j \rangle, w \rangle)] \neq \emptyset$$

As is the default in an attitude context, the anchor for any EMAs in the scope of that attitude will normally be understood to be anchored to the agent of the attitude, and thereby that agent's doxastic state will serve as the ground for the modal's Modal Base f . Hence, for *suppose*, the interpretation of EMAs in the complement will be anchored to the local center, $\odot^{\text{ENTERTAIN}}$.

Thus, *suppose* crucially involves a propositional attitude ENTERTAIN, and the imperative yields realization conditions that make it be truthfully realized just in case at realization time the agent entertains a counterfactually revised belief state. As illustrated by Moore's paradox, one cannot consistently entertain both p and *possibly not-p*. In describing the paradox, the generalization is usually restricted to beliefs, but the constraint is more general: If entertaining a state involves considering what follows in it, and if one is to discover anything interesting by that investigation, then the state itself must be consistent. Then in Moore's original illustration—*It's raining but I don't believe that it's raining*, the first conjunct requires that p be true in all the worlds in the entertained state—here the way the speaker purportedly takes the world to be, while the second imposes the impossible-to-realize requirement that *not-p* be true in at least one of them—that that belief state is still open to the possibility that *not-p* is true. Now recall that epistemic *might* presupposes that the suppositional extension of the agent's doxastic state to make the prejacent true *is consistent with that doxastic state*. For (40b), *Suppose it is raining but it might not be*, take the complement p to be of the form q and $\diamond\neg q$. Then to obtain the state to be entertained, one must apply REVDox to p , revising the addressee's belief state at realization time to yield a revised state in which both q is true in the resulting state, and $\diamond\neg q$ is also true in that state. For the second to be the case, it must be possible for $\odot^{\text{ENTERTAIN}}$ to extend her beliefs consistently to admit the possibility of $\neg q$. But since the first conjunct makes her (counterfactual) belief state one in which q is true in all the worlds in that state, the required extension with $\neg q$ is inconsistent, just as in Moore's original case. One cannot reasonably simultaneously entertain both p and *possibly not-p*. Then the counterfactual doxastic state that $\odot^{\text{ENTERTAIN}}$ is directed to entertain is inconsistent, and hence, (40b) is an *anomalous instruction*, since it asks something unreasonable of the addressee, a direction that cannot be realized.

But (40a) is quite different: There one is asked to imagine that one's counterfactual *counterparts* in the entertained counterfactual state—the derived centers—fail to believe something that's true in those worlds. We know well that our beliefs often fail to accord with the way things are, so this is an unproblematic state to entertain. It is the counterfactual belief state of $\odot^{\text{ENTERTAIN}}$ which entails that p (*it is raining*), not that of the derived centers. So no contradiction arises.

To work this out in full formal detail would require a dynamic semantics of belief revision. However, I think the sketch above provides sufficient detail to make it plausible that the semantics on offer for *suppose* and the pragmatics of the doxastic anchoring of *might* together offer a satisfying account of how Yalcin's puzzle can be resolved. The account is simple and independently motivated by what we have already said about the EMAs themselves, and by the semantics of *suppose*. And *pace* Yalcin (2007, 2010) and Stalnaker (2014:139ff), an account of this phenomenon requires no expressivist treatment of modality or special posterior context.

4.4. Scope and pseudo-scope

Though it is often claimed that epistemic modals cannot take narrow scope relative to other operators (Palmer 1990, Brennan (1993), Hacquard 2013), von Fintel & Gillies (2007b) offer the following examples to argue that this is not always the case:

- (41) Bill thinks that there might have been a mistake. [Attitude predicate over modal]
(42) Where might you have put the keys? [Question over modal]
(43) The keys might have been in the drawer. [Past over modal]
present perfect: shifts the time of the center to one earlier than ST
(44) There can't have been a mistake. [Negation over modal]
NPI epistemic *can't*

The semantics proposed above offer a different perspective on this matter. It's true that non-intensional operators tend not to scope over epistemic modals, leading to von Fintel & Iatridou's (2003) Epistemic Containment Principle:

The Epistemic Containment Principle: Epistemic modals tend to take obligatory wide scope with respect to a wide class of quantifiers.

In support of this, they offer (45):

- (45) Every candidate might win. [von Fintel & Iatridou 2003]

They claim that on the epistemic understanding of *might*, (45) "has no true reading if there is at most one winner of the election, even if there is no candidate that we know is going to lose." There *is* a possible reading of (45) with the universal wide, but the modal in this reading is dynamic, not epistemic: 'each candidate is capable of winning'. I agree that there's a lot of interference in this example from the dynamic interpretation of the modal, but perhaps that's partly because the dynamic reading is so very closely related in this case to a speaker-anchored epistemic.

But to control for interference from that possibility we can select a predicate that doesn't lend itself to being something someone is 'capable of', i.e. isn't a naturally dynamic complement. I also use a non-universal quantifier:

- (46) [spoken by a teacher whose students are having trouble with standardized tests:]
Given what I know about these kids, many of them might be amenable to working with a tutor.

I readily get an interpretation for (46) which can be paraphrased as 'there are many x s.t. x might be amenable to working with a tutor', and in fact find that more natural than the 'it might be that many are amenable' interpretation. So the epistemic containment principle is not particularly robust.

Also, as we might expect on the present account, it's very easy to get the wide scope reading of a universally quantified NP if that NP ranges over agents of an embedding attitude predicate that takes an EMA in its complement:

- (47) Every candidate believed, on the basis of his own polls, that he might win. Some of them were right, and others wrong.

But here, of course, the agents in question act as discourse centers that satisfy the doxastic anchoring presupposition of *might*. And, as with presupposition satisfaction generally, if some operator takes wider scope than the anchor for a modal, that operator will take wide scope over the modal as well:

- (48) Every year, most candidates believed, on the basis of their own polls, that they might win.

with scope order: *every* - *most* - *might*

In other words, one key to the tendency to wide scope of EMAs is actually the **pseudoscope** typical of constituents whose lexical content involves an anaphoric presupposition. For example, consider the following cases involving anaphoric Bridging:

- (49) If you park a car on a steep hill, engage the emergency brake.
(50) The gears on this car tend to slip. If you park on a hill, engage the emergency brake.

In (49), the antecedent of the anaphoric description *the emergency brake* is understood to be the (weakly familiar, Roberts 2004) emergency brake of the car introduced in the *if*-clause. Since the antecedent is understood to be non-specific—the arbitrary car parked on a steep hill, and hence falls under the scope of the operator associated with the bare conditional (per Kratzer, a universal modal operator), the definite description seems to take narrow scope relative to that operator, as well. But in (50), the weakly familiar antecedent is the emergency brake of the car referred to with the demonstrative *this car*. Since that antecedent is introduced outside of the conditional—we're talking about a particular emergency brake, a singular entity—the particularity extends to the denotation of the emergency brake itself. Thereby, though the definite description doesn't get displaced in Logical Form, it *seems* to have wide scope over the conditional operator just because its antecedent does. This is pseudoscope.²⁴

Similarly, modal anchoring for EMAs as described here is most often to one of the interlocutors, as reflected in the initial plausibility of DeRose's (1991) Speaker Inclusion Constraint. This is like the anchoring of so-called "pure" indexicals like English *I*, *we* or *you*, as reflected in the use of global-only anchoring for indexicals in Kaplan's (1979) account. It is global anchoring which gives rise to pseudo-wide-scope, and the default tendency for global anchoring which gives rise to the initial plausibility of the claim that EMAs always take wide scope. The examples considered in this section argue that this initially plausible claim is false.

²⁴ This isn't what Kratzer (1998) means by *pseudoscope*, a term she uses to refer to specific indefinites when the operator associated with the definite has its domain restricted to a singleton set, as in Schwarzschild (1999). I borrow the term because I take specificity itself to be presuppositional, the purported domain restriction an epiphenomenon arising from pseudoscope as I describe it above.

Summarizing: Epistemic Containment is more a tendency than a “principle”, and it reflects the characteristics of default anchoring and projection: pseudo-scope.

4.5 **Response patterns, agreement and what’s at-issue**²⁵

A number of important debates in the recent literature on EMAS hinge on how we understand disagreements over assertions containing EMAS, and on a variety of attested response patterns to such assertions. In this section, we take up first what it means to agree or disagree with such assertions, arguing that this has generally been misunderstood. Then we turn to consider how the response patterns can be understood and even predicted for particular contexts, in view of an independently motivated account of how what’s at-issue in a given context helps to both drive and constrain interpretation. In the course of this exploration, we motivate a revision of the notion of RELEVANCE (§2), which makes it sensitive not only to the simple truth or falsity of a target proposition, but its possibility or probability as an answer to the QUD.

4.5.1 **Faultless disagreement**

A central argument in favor of Relativism over Contextualism—in modals as in predicates of personal taste (Lasersohn 2005; Egan et al. 2005; MacFarlane 2005, 2011; Egan 2007; Stephenson 2007, etc.) has been the observation of so-called faultless disagreement. A relativist would argue that in examples like (51) neither Agnes nor Bruce has warrant to claim that what the other said was false, since whether or not something is tasty is a matter of personal taste, not subject to judgment by others:

- (51) Agnes: This salty Dutch licorice is tasty.
Bruce: No it isn’t! It’s terrible.

Now consider a modal counterpart in (52), from Egan (2007):

- (52) [Context: James Bond has just returned to London after infiltrating SPECTRE’s secret base in the Swiss Alps, planting a bug in the main conference room and slipping out by night, leaving persuasive but misleading evidence of his presence in Zurich. . . .While monitoring the newly placed bug, Bond and his CIA colleague Felix Leiter overhear the following conversation between Blofeld and his second in command, Number 2, after Number 2 has discovered the misleading evidence:

[Number 2 to Blofeld:] Bond might be in Zurich.

Upon hearing this, Leiter turns to Bond and says:

²⁵ This section owes a great deal to my work with David Beaver, Mandy Simons and Judith Tonhauser, cited throughout. However, they are not responsible for any errors in my explanation and application of our results to the cases discussed here. My understanding of Relativism owes a great deal to discussions with Kevin Scharp, Stewart Shapiro, and, especially, Eric Snyder. But again, they’re not to blame for any ways in which this account falls short.

[Leiter to Bond:] That's false.

In this scenario, both Leiter and Bond have reason to know of the misleading evidence on which Number 2's claim is based, and to know that *based on that evidence* it is possible that Bond is in Zurich. Nonetheless, according to the Relativists like Egan, Leiter's *That's false* is warranted.²⁶ They then claim that the only way to explain this is to assume that Number 2 and Leiter have a faultless disagreement about the truth of the same proposition, 'Bond might be in Zurich', their assessments of its truth differing according to their differing evidence.

In large part on the basis of such examples (see also Granger's puzzle, above), Relativists claim that the compositional, truth conditional content of a clause yields a proposition which is neither true nor false by itself. Instead, it can only be judged by appeal to a contextually given assessor, or Judge: in the case of (51), an agent whose taste determines the standards by which the licorice is judged; in EMA statements like (52), the agent whose evidence or epistemic point of view is brought to bear. Thus, in (52) Number 2 and Leiter express different opinions about the same proposition, 'Bond might be in Zurich', the former asserting it, the latter denying its truth. The assessor is a special type of index of interpretation, given once and for all for a given utterance.²⁷ In contrast, on the present account each EMA presupposes an anchor, which is contextually resolved like other presuppositions in the course of dynamic interpretation. This type of account leaves open the possibility that multiple EMAs (or other epistemic operators) within a clause might be resolved to different anchors. And it makes *Bond might be in Zurich* with *might* anchored by Number 2 express a proposition which itself takes Number 2's doxastic perspective into account. That doesn't seem to be the proposition that Leiter is denying, given that he knows about the misleading evidence.

Of course, different evidence may entail different open possibilities, without any of the relevant agents being "wrong". As a clear instance of this, consider this example from von Fintel & Gillies (2008), after Gibbard (1981):

- (53) 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:
- 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.

²⁶ I don't find Leiter's response natural in this constructed example. It would be more natural if we thought Leiter didn't know that misleading evidence was planted. Or in the scenario described, he might say *But we, of course, know better*. But since others have accepted the example as given, I'll ignore this quibble here.

²⁷ See Lasersohn (2008) for treatment of cases where judges seem to be quantificationally bound with a special meta-operator that shifts indices of assessment, and Snyder (2013) for a critique of Lasersohn's proposal.

Though this example seems to involve true faultless disagreement, von Fintel & Gillies point out that if the Boss is the assessor, or judge, in a relativist semantics, and has no information about who the turncoat is, then both Jack's and Zack's reports are false at that index, so he shouldn't draw his conclusion. But clearly he has good reason to think that both are correct, and accordingly draws the correct conclusion. So this kind of case seems to argue that the Relativist strategy for cases like (52) is inadequate to explain the full range of relevant intuitions.

Also, in (33), repeated from above, different bodies of evidence to which the same (speaker) anchor is privy lead to faultless disagreement between one agent and herself:

- (33) Given the results of the DNA tests, John might be the thief. But if we take the eyewitness seriously, John can't have been the thief. [von Fintel & Gillies 2007b]

Simply taking an assessor to be an individual agent would not suffice to explain the difference here. Perhaps the specification of the assessor could be enriched to relativize it to an agent plus a given body of information. But Dowell (2011) considers a wide range of additional examples involving (dis)agreement, and concludes that altogether they pose significant problems for Relativism, both "solipsistic Relativism", where only the assessor's evidence is relevant, and a more flexible Relativism which admits of, e.g., group evidence.

A number of researchers, including von Fintel & Gillies (2008), Dowell (2011) and Moltmann (2012) criticize Relativism and defend Contextualism for EMAS, partly based on consideration of such examples, as well as on additional truth conditional grounds. Altogether, the weight of the evidence against the Relativist approach seems compelling.²⁸ And the approach proposed here, with variable anchoring, offers us the flexibility to account for examples like (53) and (33). But there is still the question of how to give a satisfying general account of examples involving apparently faultless disagreement.

Consider a scenario in which two agents directly disagree about a modal claim:

- (54) Beau: John might be the murderer.
Cecile: No, he can't be!

Just as in (51) where Agnes and Bruce bring different individual preferences to bear in their judgments, in (54) Beau and Cecile presumably bring different evidence to bear as well. So both Relativism and the present variable anchoring account could be said to predict that their disagreement is "faultless"—Relativism because it takes the proposition *John might be the murderer* to be the same in the two assertions, just assessed differently by the two interlocutors; the variable anchoring account because it claims that *John might be the murderer* anchored to Beau expresses a different proposition than it does when anchored to Cecile, so that both utterances may express true propositions. But then, on the variable anchoring account, wherein lies the disagreement indicated by Cecile's *No*? Is it a disagreement over the truth of the modal statement itself, over the truth of the prejacent, or, as I will argue, something else?

²⁸ Again, see the previous footnote for the more general problem with quantificational variance over judges. See also Glanzberg (2007), Saebø (2009) and Schaffer (2011) for other critical discussions of Relativism.

This is related to a problem with the account of Yalcin (2007), pointed out by Stalnaker (2014:143,fn.12). Yalcin accounts for shifting epistemic anchoring via an indexed information state. He claims that this state is obligatorily shifted under attitudes to the information state of the agent of the attitude. Stalnaker takes issue with this obligatory shift, giving the following example:

(55) Alice: Jones believes that it might rain, but I disagree—I think we can count on fine weather all day.

If the information state relative to which *might* in the complement of *believes* is interpreted is that of Jones, then by asserting the first conjunct Alice is committed to the truth of *it might rain* in Jones' information state. But (55) strikes us as consistent, so we must conclude that Alice's disagreement is not over the proposition expressed by *it might rain* so-anchored, but over what possibilities are live options.

This example points up a respect in which the debate over faultless disagreement as it has been framed to this point begs a deeper question: What does it mean to agree or disagree in response to what someone has said or with what someone believes?

Moltmann (2012) talks about both faultless disagreement and what she calls "faulty agreement", illustrated by the following:

(56) Mary believes that it may rain (because she heard the weather forecast).
John believes that it may rain (because he noticed the cloud formation).
John and Mary believe the same thing (but for different reasons). [valid]

(57) Mary believes that it may rain (because she heard the weather forecast).
John believes that it may rain (because he noticed the cloud formation).
John and Mary believe different things. . . [invalid]

In (56) it's clear that there's a sense in which both Mary and John believe that *it may rain* is true, even though they don't share the same grounds for believing it. Even stronger, it's quite odd in that case to say that they believe *different* things (57). But on the present Contextualist account (as well as those of the other authors cited above), the surface-identical complements of *believes* in the two premises in these arguments express different propositions, the first about what follows from Mary's beliefs, the second about what follows from John's. So what is the sense in which the two agents agree?

This paradigm and (55) strongly suggest that (dis)agreement isn't about the truth of the proposition expressed by the speaker in a speech act or attitude complement, but about a feature of the belief states of the relevant parties. That is, what the premises in (56) and (57) and the complement in (55) tell us is that it is compatible with Mary's and John's and Jones' belief states that it will rain. And this is the sense in which they all agree: all have belief states that admit of the (future) possibility of rain. So arguably the shared belief, the agreement, is not about what the *speaker says* so much as about the possibility that it points to. Similarly, in (55) Alice disagrees with Jones in that while his belief state admits of that possibility, hers does not. This is

what agreement and disagreement are about: belief states that are congruent in the relevant way. This cannot be reduced in the general case to shared assessments of the simple truth of a proposition.

To clarify: von Fintel & Gillies (2008) argue that what Leiter is denying in (52) is not the modal claim by Number 2, but its prejacent. That seems intuitively correct for (52), and MacFarlane (2011:147) accepts that assessment of the example. However, we should not over-generalize from this to assume that all cases of (dis)agreement over EMA statements are reducible to disagreements over the truth of the prejacent in those statements. For example, in (54), Beau hasn't asserted that John is the murderer, so it seems incorrect to claim that Cecile is arguing with him about *whether John is the murderer*, i.e. the truth of the prejacent.²⁹ Nor is it intuitively clear that Beau and Cecile would agree with the theorist who claims that their disagreement is "faultless" in the sense that each is correct relative to his or her own (differing) evidence. Rather, even if they explicitly relativize their claims to their own belief states—*Given what I know, John might be the murderer./No! Given what I know, he can't be!*—they would say that one of them is right and the other is wrong, and that they're arguing about whether John's being the murderer is within the realm of possibility. That is, their disagreement is not about the truth of a proposition (either the prejacent or the evidential claim) but about what a realistic belief state should be like. Something similar can be said about (55) and (56)-(57), as well. Belief states are richer than propositions, and they admit of a richer notion of agreement and disagreement: about possibility and necessity (and likelihood, etc.) as well as simple truth or falsity.

This shines light on some of the earlier examples considered, as well. In (52), Leiter disputes not so much what Number 2 has said as what Blofield thereby admits as possible, i.e. Bond being in Zurich. In (53), in accepting what both his trusted informants tell him, the Boss comes to have a belief state that accords in the relevant ways with *both* of theirs; but he thereby has more information than either of them and can reach a firm conclusion about who the turncoat must be.

Hence, the recent debate over faultless disagreement in utterances involving EMAs seems to be somewhat misguided. One can argue that retractions and disavowals are similarly about one's former beliefs, rather than necessarily about the truth of the statements that reported them. And one might speculate that this is more generally the case, pertaining to the debate over predicates of personal taste as well.

But wait a minute! It's fine to say that we have opinions about what a belief state should be like with respect to propositions like those denoted by the above prejacent (making them true, false, possible, probable, etc.). But don't we also have beliefs about others' beliefs? And couldn't we disagree about *those* beliefs, too? As we'll soon see, we do just that in cases involving group investigations based on joint evidence and in games like Mastermind.

²⁹ One might argue that in (202) Cecile has just asserted that the prejacent is false, so that her disagreement with Beau is over its truth value. But this is not generally sufficient: We could change the discourse to:

- (i) Beau: I think it's likely that John is the murderer.
Cecile: No, that's very unlikely/implausible.

Then though Cecile seems to concede that John being the murderer is in the realm of possibility, she disagrees with Beau about its likelihood. Again, the disagreement lies in the shape of their belief states, not in their assessment of the truth of a simple, non-modal prejacent or complement. See Moss (2015) on graded epistemic modals like *likely*.

But if that is so, how do we know, in a particular case, just what the disagreement is about?

In the next section, I'll argue that a (dis)agreement in a particular discourse, and in other types of judgment and response as well, are not necessarily directed at the entire proposition expressed by the target utterance, but at what's at-issue in that utterance. Since QUD theory tells us that interlocutors carefully track what's at-issue at any point in a discourse, this gives them adequate grounds to understand the disagreements expressed.

4.5.2 Responding to what's at-issue

Patterns of felicitous evaluative response to statements containing EMAS have been the subject of extensive discussion in the literature. See especially Lyons (1977), Swanson (2006), Stephenson (2007), von Stechow & Gillies (2007b,2008), Portner (2009), and Dowell (2011), among many others. Response patterns have been used as evidence for a variety of proposals about the semantics of EMAS, ranging from speech-act accounts of modality through modal Relativism to cloud-of-contexts accounts. But I am not aware of a systematic consideration of the range of possible evaluative responses and their import for the semantics of EMAS. Here I characterize one general feature of what response patterns tell us about the discourse status of an utterance: They reflect what QUD the responder takes the target utterance to address, i.e. what's at-issue. And interpretation of the response itself requires the same information. And I consider what that might tell us about the status of the prejacent in such utterances vs. that of the evidential claim itself.

Consider a simple discourse involving responses to statements with EMAS:

(58) QUD: What was the weather like yesterday when we were out of town?

Alex: It might have rained yesterday.

Possible responses

Barbara: No, it didn't.

Chris: No, it can't have rained. I would have seen it on the weather report.

Assume that *might* in (58A) is anchored to Alex's belief state, or to the Common Ground as he knows it. The tag on Barbara's response makes it clear that what's being disputed is the truth of the prejacent, whereas the alternative response by Chris seems to target the possibility that it was raining. Of course, a rejection of the prejacent would entail a rejection of the possibility that it rained; but on most accounts of epistemic *might*, the present account included, *might p* doesn't proffer the *mere possibility* that *p*, but instead its possibility relative to the worlds in the anchor Alex's belief state. However, as we saw in discussing purported faultless disagreements, our intuition tells us that this relativized possibility is not what Barbara is getting at in (58)—in fact, it's not what Chris seems to be rejecting either.

Taking into account related observations about felicitous responses to declarative statements with EMAS in the root clause, Swanson (2006. §2.3:73ff) argues that on some occasions someone saying *might p*, serves to performatively raise the possibility that *p*, while on others it may be used to either assert something (proposing to add the proposition to the CG), or to advise an

addressee on how to update her subjective probabilities. That seems intuitively correct from a functional point of view. But this only raises further questions: (a) Do we need three kinds of content for *might p* to explain how it can serve these different kinds of functions?, and (b) Are there systematic contextual constraints on which function we can take *might p* to serve in a particular utterance? If the answer to (b) is *yes*, as I will argue, then we might consider whether the different roles *might p* plays are a function of the context of utterance, reflecting the pragmatics of modal talk, rather than motivating some of the special semantics argued for by other authors (a).

In connection with question (b), it's clear from context in some examples that the doxastic state associated with the EMA is part of what's addressed in the response:

- (59) [Watson to Holmes, both privy to the same body of evidence:] Then the butler must be the murderer!
[Homes to Watson, patiently:] No, my dear Watson. You're forgetting that the housekeeper also had the keys to the wine cellar.

Here Holmes doesn't mean that the butler *isn't* the murderer, but instead refutes Watson's claim about what follows from the evidence available to them at that point—an inclusive 'we' (©^{CG}) anchoring of the EMA. This disagreement is about what the interlocutors' joint evidence entails, and accordingly about how their belief states should agree.

A set of related examples considered by von Fintel & Gillies (2008:83-84) consist of (constructed) dialogues between players of the game Mastermind. This game involves a pair of players, one of whom, the codemaker, has full access to an array of colored pegs, the other trying to guess the colors and pattern within a limited period. This usually involves a series of partial guesses by the codebreaker about the locations of particular colors. At each turn, the knowledgeable codemaker can give positive feedback about what the codebreaker gets right. Here is one of von Fintel & Gillies' examples:

- (60) [Pascal and Mordecai are playing Mastermind. After some rounds where Mordecai gives Pascal hints about the solution, Pascal says:] There might be two reds.
- (61) [Mordecai, knowing the solution, has a range of possible responses:]
- 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.

Mordecai's replies (61a) and (61c) clearly respond to the evidential claim in (60)—they affirm or deny whether Pascal correctly reported his own evidential state, while (61b) and (61d) respond to the prejacent alone. If Mordecai is playing his role properly—rather like Holmes' pedagogical role in relation to Watson, we expect a response that gives Pascal feedback about his evidence (a or c) instead of announcing the correct solution (b or d). Note also that response (61c) would express Mordecai's disagreement with Pascal about what Pascal's own belief state should be (on the basis of the evidence Mordecai knows Pascal has had access to), rather than expressing

disagreement about the simple prejacent alone. This is a solipsistic anchoring of *can't*, with the same anchor as that of *might* in (60).

But what can we say about EMA evaluative response patterns more generally and about *where we might expect to find which responses*? The hypothesis I will consider here is that patterns of felicitous evaluative response reflect what's at-issue in the utterance to which they respond, in the technical sense of *at-issueness* defined in §2.³⁰ In support of this, I will bring to bear several diagnostic patterns which give evidence about what's at-issue in an utterance, i.e. about what makes it RELEVANT to the QUD in the context of utterance (Roberts et al. 2009, Simons et al. 2011, Tonhauser 2012).³¹ First, in §4.5.2.1 we consider independence evidence for their utility, and then apply them to EMA statements and questions. Then in §4.5.2.2 we briefly consider how this bears on RELEVANCE, arguing for a refinement of that notion to take into account talk of possibilities.

4.5.2.1 Apt response, RELEVANCE, and diagnostics for what's at-issue

Recall the different types of content that may occur in the Character of an expression—presupposed content, auxiliary content, and proffered content. Non-proffered content cannot satisfy the requirement that an utterance be RELEVANT to the QUD. For example, Potts (2005) claims that CIs (which crucially contain auxiliary content) are “nondeniable”, and Amaral et al. (2007) extend this observation to demonstrate that we normally cannot directly respond to an appositive like that in (62) with a direct affirmation or denial. In this example, the denial can only be understood to target whether Edna started the descent, and not whether she's fearless. To respond to the appositive itself, we need something like *Hey! Wait a minute!*, as in (63):

(62) A: Edna, a fearless leader, started the descent.
B: No, that's not true.

(63) A: Edna, a fearless leader, started the descent.
B: Hey, wait a minute! – Edna is not a fearless leader. She's a coward!

These illustrate the following principle, Apt Response:

Apt Response: A direct response to utterance *U* can only be understood to target the at-issue content of *U*.

A direct response might be evaluative: e.g., *yes, no, that's true/false*, or in case *U* is a question, it is a (possibly partial) answer to *U*. Apt Response has implications both for the interpretation of

³⁰ The structure of discourse is quite complex, and accordingly, responses can serve a variety of functions. E.g., see Asher & Lascarides (2003), Ginzberg (2011), and Farkas & Bruce (2010) for other important considerations. Hence, this hypothesis is *not* intended to serve as a complete theory of the semantics and pragmatics of response.

Instead, I focus here on one feature of responses that's independently motivated and relevant to the subject at hand.

³¹ The inventory of ways of responding directly to an utterance, and what they presuppose and implicate about the response, vary across languages: see Farkas & Bruce (2010) for very interesting discussion of ways of confirming or denying and how they may differ from language to language. Here I simply rely on native intuitions of speakers of English as they pertain to the meaning of the English responses considered.

the target utterance and for that of the response itself. Insofar as the response is clear, it can shed light on what's at-issue in the utterance that provoked it; and the response itself is interpreted in light of what's evidently at-issue in the provoking utterance.

Recall the definitions of RELEVANCE and of what it is to be at-issue from §2:

- RELEVANCE to the QUD:** Felicity of utterance requires RELEVANCE to the QUD, and
- An assertion is relevant to a QUD iff it contextually entails a partial or complete answer to the QUD.
 - A question is relevant to a QUD iff it has an answer which contextually entails a partial or complete answer to the QUD.
 - A directive is RELEVANT to a QUD iff its realization promises to play a role in resolving the QUD

at-issueness: For any proposition p , let $?p$ denote the question whether p , i.e. the partition on the set of worlds with members p and $\neg p$. Then:
A proffered proposition p is **at-issue** relative to a question Q iff $?p$ is RELEVANT to Q .

Take a Provocation/Response sequence in discourse to consist of a speech act (question, assertion, direction)—the *Provocation*, that prompts a *Response*. Together, Apt Response, the requirement of RELEVANCE and the above definitions interact to yield a set of predictions about what constitutes a felicitous sequence of Provocation and direct Response, and about how these will be interpreted, yielding the constraints summarized in Table 2, where (N)AI is an abbreviation for (Not) At-Issue:

Character of Interaction		Constraints as a function of (Not) At-Issue content		
Provocation	Response	AI content of Provocation	AI content of Response	Constrains
Question	Answer	Question Relevance: the NAI content of a Question doesn't bear on an Answer's Relevance	Answer Relevance: the NAI content of an Answer does not satisfy Relevance to a Question	Relevance to Question
Assertion	Evaluation	Assertion Relevance: the NAI content of an Assertion is not the target of an Evaluation	Evaluation Relevance: the NAI content of an Evaluation needn't play a role in evaluating the Assertion	Nature of Evaluation
Constraint follows from:		Apt Response	def'ns of <i>RELEVANCE</i> , <i>at-issue</i>	

Table 2: **Constraints on Provocation/Response interpretation**

In the third column, Apt Response generalizations Question Relevance and Assertion Relevance are instantiations of Apt Response, constraining the range of felicitous responses to a Provocation in a given context, and predicting which aspect(s) of the content of the Provocation a direct response would target. When the Provocation is a question, Question Relevance tells us that it is only the at-issue content of the question that determines what it is to be a RELEVANT

answer. (62) and (63) above illustrate Assertion Relevance:³² only the AI content in A's utterance in (62A)/(63A) can be the target of the direct denial in (62B), the NAI CI content requiring the non-direct *Hey! Wait a minute!* response. Similarly, according to Assertion Relevance, when the Provocation is an assertion, only its at-issue content will determine what counts as the felicitous target of a direct evaluative response. Question Relevance is illustrated in (64), wherein the yes/no question itself has both at-issue content and not-at-issue CI content:

- (64) A: Has Edna, a fearless leader, started the descent?
B: Yes (she has started)./No (she hasn't started).
B': #Yes, she's fearless./#No, she's not fearless.

The question is whether Edna has started the descent, not whether Edna is a fearless leader (the CI content), and what counts as a felicitous direct response differs accordingly.

Among the RELEVANCE-based patterns in the fourth column of Table 2, Answer Relevance is illustrated for CI content by (65).³³ A *wh*-question cannot be felicitously addressed by an utterance in which the content of a CI is crucial to ensure the response's RELEVANCE to the question:

- (65) Who's Edward Witten?
a. #A former linguist, Edward Witten, is now the top-dog in string theory.
b. Edward Witten, a former linguist, is now the top-dog in string theory.
(Amaral, Roberts & Smith 2007)

In (65a), the main clause without the appositive, *a former linguist is now the top-dog in string theory*, does not by itself address the QUD posed by the explicit question *Who's Edward Witten?*. Since only at-issue content can satisfy the requirement of RELEVANCE, and the RELEVANT appositive is not at-issue, (65a) is infelicitous. In (65b), the contents of subject and appositive are switched, yielding a RELEVANT at-issue main clause, and the result is felicitous.

Finally, Evaluation Relevance is illustrated for CIs by (66), where (66B) is a very odd response to (66A) because the only content in (B) that's relevant to (A), justifying (B)'s rejection, is the auxiliary content of a CI:

- (66) A: Edward Witten is from St. Louis.
B: #No, Witten—who is from Baltimore, not St. Louis—is a physicist working on string theory.

The results from application of all of these diagnostics argue that appositives are normally not at-issue.³⁴ Hence, they are not part of the proffered content of the utterance in which they occur. But the diagnostics work not only for the auxiliary content of CIs, but for other types of not at-

³² See Tonhauser's (2012) carefully constructed diagnostics (1a, 1b, 1c), designed for using Question Relevance to explore what's at-issue in fieldwork elicitation or through controlled experiment.

³³ See Tonhauser's (2012) Diagnostic 2.

³⁴ There are at-issue uses of appositives and non-restrictive relative clauses. See Syrett & Koev (to appear). But those occur under special circumstances and will not concern us here.

issue content as well, including content which is locally entailed—contributing to the truth conditional content of the clause in which it occurs, and hence proffered.

For example, factive verbs like *know*, *discover*, and *be aware of* are standardly said to presuppose the truth of their complements, as well as entailing them. Hence, on the standard accounts of presupposition (Karttunen 1973, Heim 1983), we would expect that they always give rise to presupposition projection, wherein their complements are taken to be true even when they fall under the scope of negation, interrogation or a modal, despite the fact that these operators are normally entailment cancelling. But it is well-known that in some cases factive complements fail to project in that way; Abusch (2002,2009) thus dubs them “soft presupposition triggers”. Beaver (2010) offers numerous naturally occurring examples when projection fails to arise with these triggers. Simons et al. (2011) and Simons et al. (2015) argue that in fact projection only occurs when the truth of the complement is not at-issue, whereas it fails to occur when the complement is at-issue, whether with respect to an explicit interrogative or, as in the following naturally occurring example, due to an implicit QUD:

(67) ...I haven't tried this with wombats though, and if anyone discovers that the method is also wombat-proof, I'd really like to know. (Beaver 2010, (32))

In (67), the first conjunct entails that so far as the speaker knows the method (*this*) may not work with wombats, and the remainder makes it clear that the speaker would like to know whether it does. So the complement of factive *discovers*, *the method is also wombat-proof*, is at-issue. Accordingly, it fails to project out of the antecedent of the conditional, a typical projection environment. See Beaver for many other examples, Simons et al. (2015) for further discussion. And Roberts (2011) for examples where the usually projective prejacent of *only* also fails to project when it's at-issue.

This is just to say that when the complement of a factive or the prejacent of *only* is that portion of the utterance which satisfies RELEVANCE, its truth cannot felicitously be taken for granted in the same context of utterance—it cannot be presupposed in the sense of Stalnaker (1974)—since that would amount to a pragmatic contradiction: if a QUD is felicitous, the CG cannot entail its complete answer (Roberts 1996). Therefore, in such contexts the complement does not project, a behavior only licensed by pragmatic presupposition. Besides the diagnostic patterns discussed here, there is a growing body of experimental evidence that this is the case with the factives (Xue & Onea 2011; Cummins, Amaral & Katsos 2012; Koev 2013; Syrett & Koev to appear), and with the normally factive evaluative adjectives as well (Tonhauser et al. 2015).

This failure of factive predicates to trigger presupposition when the complement is not at-issue is closely related to a phenomenon discussed by Simons (2006) (drawing on observations due to Hooper 1975), wherein a class of verbs that take sentential complements (e.g. *see*, *hear*, *think*, *believe*, *discover*) may behave instead as parenthetical. On this use, “the embedded clause carries the main point of the utterance, while the main clause serves some discourse function”. Here are some examples:

(68) Why hasn't Louise been coming to our meetings recently?
a. I believe she's left town.
b. She's left town, I believe.

- c: Henry thinks that she's left town.
Possible replies to a - c:
d. But she hasn't. I saw her yesterday in the supermarket.
e. No he doesn't. He told me her saw her yesterday in the supermarket.

Simons argues that the parenthetical uses of *believe* and *think* in (68a-c) are evidential. This type of use tends to license slifting, as we see in (68b), wherein the subject and main predicate are postposed, rather like an adverbial. In (68a-c), what's at-issue relative to the explicit QUD isn't the speaker's or Henry's doxastic state, but the proposition denoted by the complement, 'Louise has left town'. And this is reflected in the possibility of a direct response to the truth of that complement, (68d).³⁵ But even though this use is parenthetical, that does not mean that the attitude predicate plays no role in the proposition expressed: the speaker in uttering one of (68a) – (68c) is not committed to the truth of the complement *per se*. As with EMAS, where the speaker can deny commitment to the truth of the prejacent, the evidential in one of these assertions amounts to a hedge. And accordingly, someone can respond to one of them by denying the evidential statement instead of the prejacent itself, as in (68e), thereby weakening the CG evidence for the truth of the RELEVANT prejacent.

von Stechow & Gillies (2007b) notice that the parenthetical behavior of predicates like those in (68), is parallel to that of some uses of EMAS, as in their (69):³⁶

- (69) Q: Why isn't Louise coming to our meetings these days?
A: She might/must be too busy with her dissertation.

In this context, only the prejacent of (69A) is directly RELEVANT to the explicit question. Thus Answer Relevance argues that only the complements are at-issue in (68), the prejacent in (69A). I.e., the evidential content of the EMA is pragmatically backgrounded in (69A). Just as (68a)-(68c) fail to entail their complements, the proffered content of (69A) does not entail the truth of the prejacent. The speaker is not committed to the truth of the foreground, at-issue

³⁵ Further evidence that in parenthetically used attitude reports the complement is at-issue comes from a subclass of the parentheticals, verbs that Karttunen (1971) called the *semi-factives* (including *found out*, *realize*, *learn*, and *discover*, among others), so-called because they have a strong tendency to presuppose their complements but quite often (even in non-parenthetical uses) do not do so. But when are they presuppositional? As Simons notes, the standard examples in the literature are "cases where the expected presupposition disappears along with the factive implication", examples like (67) above. The semi-factive *discover* occurs in that example in a context where presuppositions typically project, in the antecedent of a conditional, but there is no projection—no implication of the truth of the complement. In that case, as we saw, the presupposition would be incompatible with the speaker's implication that the truth of the prejacent is at-issue.

The parenthetical use of a semi-factive, in contrast, is unembedded and therefore entails the truth of its complement. But instead of being presupposed, the complement is treated as new information, the main point of the assertion, as in (i), with a set of parenthetical semi-factives that could be felicitously used here:

- (i) Q: Why isn't Louise coming to our meetings these days?
A: Henry's discovered / realized / figured out / learned that she left town. [after Simons 2006]

What the two kinds of examples, (67) and (iA), share is that the complement is not at-issue, made clear by the explicit QUD (iQ). But here, as with the non-factive parentheticals in (68), we can also directly respond to the prejacent with (68d). Thus, the thesis that only not at-issue content projects, evident in (67), converges here with evidence from Apt Response and Relevance that the complements of parenthetical semi-factives are not at-issue.

³⁶ They do not analyze the example in terms of what's at-issue and Relevance, but I take it that what they say is compatible with the present explanation.

complements/prejacent, only to their RELEVANCE. Then Apt Response (Assertion Relevance) leads us to correctly predict that responses may target the prejacent in such uses, as in (69R), a perfectly felicitous direct response to (69A):

(69') R: No she isn't—she filed it last month!

R': That can't be right! I saw her out partying several times in the last month.

The fact that (69'R') is also felicitous doesn't argue either that the evidentiality itself is at-issue or that *might* and *can't* are anchored to the same agent(s): As we saw above in (54), even with *can't* anchored to the speaker of R' such an indirect response entails an unwillingness to accept the prejacent as a (possible) answer to (69Q). This illustrates how the independently motivated principles and associated diagnostics considered here seem to apply naturally to EMA statements, as well.

Assertion Relevance can also straightforwardly explain the responses in (58) above, assuming the anchoring of the modals as discussed. Barbara's felicitous response (58B) targets the prejacent, taking it to be what's at-issue in Alex's assertion, even though Alex did not directly assert the prejacent. The explicit QUD *what was the weather like...?* is not itself directly addressed by information about anyone's belief state. Only the prejacent of (58A) is directly RELEVANT to the QUD, and that is what is targeted by the felicitous (58B). Moreover, Chris's own (solipsistic) rejection of the possibility that it rained (58C) entails a rejection of the prejacent itself; hence is also contextually RELEVANT to the at-issue prejacent. Another way of putting this is that Chris's response rejects Alex's proposal that they entertain the possibility that it rained. This accounts for the felicity of (58C) without taking the Chris-anchored response to directly entail the falsity of Alex's belief state as reported in (58A). The interlocutors are negotiating not about the truth of (58A), but about how to update the CG.

Prima facie, examples like (58) have been taken to make the best possible case for Relativism: Alex and Chris both make claims about the possibility that it rained, one promoting it, the other denying, and Barbara's assertion likewise entails that that's not a possibility. The fact that each of them may make these claims based on their own solipsistic evidence doesn't seem to figure in their *responses*. These are uses where, as Swanson (2006:73) puts it "someone saying *might p*, serves to performatively raise the possibility that *p*".

Such examples are quite different from the Holmes/Watson example (59) and the Mastermind responses in (61). Talk of the famous Holmes and his sidekick Watson evokes a context in which we know about the kind of thing they're up to—sifting through evidence to try to reach a best-hypothesis. The goals and intentions of rational agents constrain and, in cases like this, establish the QUD, what's at-issue, as argued in Roberts (2004,2012b). In a related note, Dowell talks about how understanding the relationship between an agent's information state and her actions can guide the interpretive resolution of a EMA's anchor (not her term): "Contexts in which a speaker is offering a [Basic Epistemic Modal assertion] as an explanation of her action, or in which another person offers one as an explanation for another's, force us to hear the relevant body of information narrowly [i.e., as information to which the actor is privy]; after all, an actor can't be intending to act on information that is beyond her powers to possess" (2011:17). All this is consistent with the standard Belief-Desire-Intention model of agency.

Since Watson and Holmes are discussing their joint evidential state, the evidentiality and force of the EMA are RELEVANT, and Assertion Relevance correctly predicts that Holmes' reply targets not the prejacent p , 'the butler is the murderer', but the entire modal proposition, the claim that the evidence entails p .

In the context of (60), since the goal of the game is for Pascal to guess the solution, and the dialogue is intended to serve that goal, any questions cooperatively posed or addressed should be RELEVANT to that goal. Roughly, an action is relevant to a goal just in case it promises to help to achieve that goal. Then the question that's RELEVANT for the goals of the Mastermind game isn't one whose answer would provide the final solution of the puzzle, but instead one whose answer gives Pascal clues to that solution; this question is something like 'Is Pascal's current doxastic state in accord with the evidence available to him at this point in the game?' In the context of the game, Assertion Relevance predicts that replying to the prejacent alone, as in (61b/d), would be uncooperative—it would both fail to address the QUD (which is about what Pascal's state of mind should be) and give the game away, depriving Pascal of the opportunity to win. The RELEVANT responses would be (61a) or (61c).

In some examples, direct denials of truth may be understood to either target the entire evidential or the prejacent alone. Consider Portner's (70):

- (70) Andy: Max must be lonely.
Bob: That's not true.

Portner takes Bob's *That's not true* to have a similar meaning to that of *I agree*, which he takes to mean 'I agree that *Max must be lonely* is a reasonable thing to say'. But I respectfully disagree. First, note that taking Andy's statement out of the blue like this leaves us with no clues about what's under discussion, except that it's probably Max's state. The only reasonable anchors would be either the speaker at utterance time ©*, or the inclusive 'we' ©^{CG}. Then if the interlocutors are talking about *how Max is*, either understanding of the response would be RELEVANT. We can imagine two kinds of follow-up to (70B): (i) *That's not true—according to Marcie he may have a secret boyfriend* or (ii) *That's not true—though he's a loner at work, he's found a great community in his church*. Hence, in such a context, *That's not true* alone is indeterminate between a rejection of the necessity of Max's loneliness (whether the rejection is anchored to Bob's evidence, or to that of a third party as in (i), in either case proposing a revision to the common ground whereby it admits of the possibility of Max not being lonely) vs. a rejection of the at-issue prejacent 'Max is lonely' alone (as in (ii)).

Now let's briefly consider how EMAs work in interrogative utterances, and the role they play under Question Relevance in determining what answers are felicitous and what those answers can be taken to mean.

Consider (71):³⁷

³⁷ Inspired by an example in von Stechow & Gillies (2007a). Would that there were space to analyze their original example in detail here. See Dowell (2011) for enlightening discussion.

- (71) [Context: Alex and Billy are roommates:]
Billy: Where are my keys?
Alex: I have no idea. Where might you have left them?

Billy's question implicates that she doesn't know where her keys are. Then Alex's question seems to display interrogative flip, so that the default assumption in the context given is that *might* is solipsistically anchored to the addressee Billy. Alex is probing for any information Billy might have that would help to locate the keys. (There is another question Alex might have asked: *Let's see now—where might you have left them?*, which has more the flavor of a rhetorically posed question for joint reflection, taking the anchor of *might* to be ©^{CG}, inclusive 'we'.) Here are some answers Billy could felicitously make:

- (71') Billy: I don't know.
I thought I might have locked them in the car, but I didn't see them through the window. They're not in the kitchen. . .

But compare Alex's question in (71) and that in (72):

- (72) Alex: Where are your keys?
Billy: I don't know.
I might have locked them in the car...

Crucial here is that Billy's reply *I don't know* in the two cases expresses a different proposition. In (72), *I don't know* means 'I don't know where my keys are', whereas in (71'), in reply to Billy's *Where might you have left them?*, it means 'I don't know where I might have left them'. That is, the responses that would satisfy RELEVANCE in the two contexts are not the same. This is predicted by Question Relevance, on the assumption that the EMA is part of the at-issue content in Alex's question in (71), whereas there is no modal in her question in (72).

Finally, I claimed in §4.2 that anchoring an EMA to the addressee in an interrogative—interrogative flip—is not automatic, but is another reflection of the context of utterance. This is illustrated in the following Mastermind case:

- (73) [After some rounds where Mordecai gives Pascal hints about the solution, Pascal asks:]
{Do there have to / Must there} be two reds?

Achieving the over-arching goal of the game depends on the state of *Pascal's* evidence, not Mordecai's. In fact, Pascal knows that Mordecai knows the truth of the prejacet itself. So here the only goal-RELEVANT anchoring of *have to* or *must* is Pascal, the speaker himself. Hence, the intentional structure of the game constrains our understanding of the question. The result is a question with an EMA which does *not* involve interrogative flip. And again, the most felicitous, cooperative and direct responses to (73) are those like the following:

- (73') [Mordecai:] No.
(73'') [Mordecai:] Well, that I cannot tell you, but there *might* be.

in which *no* means not ‘no there aren’t two reds’ but ‘no there might not be two reds’. And the contrastive focus on *might* (suggested by capitalization) contrasts it with *must* in the question, arguing that the answer the speaker denies knowing in the first conjunct of (73’’) is about the necessity of their being two reds *given the information Pascal ought to control*. So RELEVANCE to the domain goal gives us the anchoring for the EMA in (73), and Answer Relevance constrains interpretation of the responses in (73’) and (73’’).

Question Relevance in interrogatives that contain EMAs shows us that the flexibility of the doxastic anchoring of the EMA is independent of its foreground/background status—the anchor may be the speaker or the addressee (flip), but in either case asking an EMA question, as opposed to questioning the prejacent alone, is likely to make the evidentiality itself part of what’s at-issue, as reflected in the interpretation of apt responses.

Here we have seen how the requirements of RELEVANCE and Apt Response, and the resulting constraints on Provocation/Response relations, especially Question Relevance, Assertion Relevance and Answer Relevance, correctly predict attested patterns of response in discourses involving EMAs. This account of the oft-attested difference between the evidential implication and the prejacent in EMA assertions is non-ad-hoc, since the phenomenon of proffered content being backgrounded by the QUD is attested in other types of content—semi-factives and several other attitude predicates, the prejacent of *only* (Roberts 2011), etc.

4.5.2.2 RELEVANCE revisited

The kinds of examples we’re considering here motivate a reconsideration of what it is to be RELEVANT, to permit it to take into account not only the simple truth or falsity of possible answers to the QUD, but their possibility or likelihood, as well. Consider (74):

- (74) Q: Is it raining?
A: It might be: People are coming in with wet umbrellas.
Well, John thinks so.
The weather report said there was a 0% chance of rain today.
They haven’t cancelled the company picnic.

None of the answers in (74A) satisfies the original definition of RELEVANCE, because none contextually entails that it is or isn’t raining. But intuitively all are RELEVANT because they bear on the *likelihood* of one or the other answer. If accepted, the CG will entail that the anchor has RELEVANT evidence. If, in addition, the interlocutors take the anchor to be reliable (e.g., John is a weatherman, the weather report is more often than not correct, etc.) and the evidence plausible, this will entail an adjustment in the CG of the probability that it is raining, and hence about whether that might be the correct answer to the QUD. And in this way, acceptance of the assertion leads to the CG agreeing with the anchoring agent’s evidential state as that bears on the truth of *it’s raining*, in the sense of *agreement* in §4.5.1.

This argues that RELEVANCE is a richer notion than is captured in the definition in §2. An assertion is RELEVANT to the QUD not only when it entails a partial answer to that question, but

also when it brings new evidence to bear on what the answer *might be*. Here is a first stab at how we might capture this for assertions:³⁸

An assertion is **RELEVANT** to the QUD just in case it contextually entails information about the likelihood of a partial answer to the QUD.

Consider cases where p directly entails a partial answer to the QUD. The speaker is always the anchoring agent in a simple assertion of p . Hence, accepting the assertion of p and coming to agree with the speaker would amount to accepting the truth of p in the CG, removing $\neg p$ worlds and thereby partially resolving the QUD. Thus in this simple case, the new notion of RELEVANCE agrees with the earlier notion in §2. Similarly, when the QUD is about some agent's evidential state and the response is an assertion EMA- p that reports on that state, the response itself is RELEVANT, and adding it to the CG partially answers the QUD.

But we can understand why the speaker might assert EMA- p even when only p itself, and not the evidential state reported, is RELEVANT to the QUD. The speaker may not have sufficient evidence to confidently assert p , but on the basis of what evidence she has thinks that p may be or even is very likely to be the answer to the QUD. Let's assume for the purposes of discussion that her evidence is purely private, not available to the other interlocutors (or that the anchor is some third person whose evidence is similarly unavailable in the CG). Then what is the result of acceptance by the interlocutors of the solipsistically anchored assertion that EMA- p ? First, as in Stalnaker (1979), it leads to the addition of the evidential proposition expressed to the CG. But in such a context acceptance also has another, crucial effect, driven by RELEVANCE: Since the interlocutors take it that the speaker knows what's at-issue, and hence that only the prejacent p is RELEVANT, and since her assertion *provides evidence about the truth of p* , it pragmatically poses the RELEVANT question $?p$. (Note that if p is RELEVANT, $?p$ is RELEVANT, and *vice versa*.) $?p$ is a partition over the CG, felicitous only if $?p$ is RELEVANT to the previous QUD and the answer to $?p$ is not entailed by the CG—i.e. both p and $\neg p$ must be live options. So by raising the RELEVANT $?p$ and thereby implying that $\not\phi p$, the speaker is thereby suggesting that the interlocutors entertain p as a possible answer to the QUD.

Just as in the parenthetical uses of attitude predicates where the complement alone is not asserted or entailed, in an EMA- p answer the prejacent is neither asserted nor entailed. This is why we say that the EMAs in such cases are proffered but backgrounded. But still what's at-issue is p , and this is reflected in what counts as an Apt Response, as we saw in the previous section.

Now consider again the three functions that Swanson (2006) suggested for an epistemic modal claim: (1) it may assert something (proposing to add the proposition to the CG); (2) it may performatively raise the possibility that p ; and (3) it may advise an addressee on how to update her subjective probabilities. I assume that in all the declarative uses considered above, the EMA- p statement is a Stalnakerian assertion, performing function (1). But when what's at-issue is p alone, (2) the statement pragmatically (via RELEVANCE) raises the possibility that p (posing $?p$), and (3) suggests a revision of the CG in such a way as to agree with the EMA-anchor, which (because the CG is *common*) entails that cooperative addressees will update their purported

³⁸ Consideration of all the factors that should be taken into account in developing an adequate revision of Relevance is outside the scope of this paper. This rough revision will have to suffice for present purposes.

subjective probabilities with respect to p accordingly. Thus, we don't need Relativism or Speech Act operators to explain the attested response patterns, nor do we need to assume that the different discourse functions performed by EMA statements reflect any conventional ambiguity.

5. Conclusions and prospects

This paper argues for three principal theses about epistemic modals like English *must* and *might*:

1. Following von Stechow & Gillies, I argue that epistemic modals are evidential. My semantics differs from theirs in significant ways, largely because I take evidentiality to be doxastic and suppositional, rather than truly epistemic and explicitly “indirect”. I claim this revision lets us capture the intuitive indirectness as a non-detachable Gricean implicature.
2. Such epistemic modals are presuppositionally anchored to a discourse center, as evidenced by (a) shiftability under attitudes, in questions, and in extended supposition, modal subordination and FID, and (b) related apparent scope patterns (presuppositional pseudo-scope). The contexts in (a) are just those in which we expect an alternative discourse center to be available and relevant. Because the default epistemic agent in an utterance is the speaker, and anchoring to the speaker gives the effect of widest scope, we predict the default wide pseudo-scope.
3. The discourse status in an epistemic modal statement of the modality *per se*, partly evidenced by attested felicitous response patterns and agreement, is a function of the QUD and what's at-issue. Epistemic modality is quite often backgrounded. However, unlike the presupposed anchoring to a discourse center, the tendency to be backgrounded is not *conventional*, but is instead a function of what's typically at-issue—the kinds of questions we tend to pose. Another way of putting this is that **we respond to an assertion *qua* answer**.

These theses are related. For example, which agent the modal is understood to be anchored to is also partly a function of the QUD, as we saw in §4.5.

Egan et al. (2005) argue against a Contextualist account of *might* and *must* partly on the grounds that their context sensitivity is different from that of purely anaphoric elements like *local* or *enemy*, which are much more promiscuous in how their implicit arguments are anchored. That is, their (75) would be felicitous without prior discussion or relevance of any particular locale:

(75) Many local bars are full of Browns fans.

Professor Granger (still in the South Pacific), hearing this spoken over Skype by her friend Jason in Cleveland, will not hesitate to say “that's right”. The fact that the relevant bars aren't local to her doesn't interfere with her willingness to agree with (75), unlike how the fact that she knew that she wasn't in Prague interfered with her willingness to agree with Myles' claim that she might be in Prague. Cross-linguistically, the difference between the EMAs and the merely anaphoric elements is real, but the reason is that the EMAs, like other indexicals—pure, impure, and shifted—require a discourse center as antecedent anchor, while *local* and its ilk do not.

Accounts which address the kinds of puzzles we've considered in §4 above by taking EMAs to be expressive, or to be interpreted relative to an Assessor (Relativism) or speech act operators, all

fail to be able to account for the full range of puzzles considered there. Furthermore, most of them have difficulties with the fact that EMAS work quite as well in embedded contexts, without many of the features that motivate special semantics in root clauses.

On the other hand, Portner (2005, 2007) and Moltmann (2012) both tentatively suggest that EMAS may be shiftable indexicals. This certainly touches on the proposal made here, but it lacks specificity, and fails to consider the difference between the shiftable indexicals, which typically can only be anchored to a lexically specified subset of the set of possible discourse centers and the EMAS, which can be anchored to any salient, RELEVANT discourse center. On the accounts I've proposed (see Roberts 2015 for the shiftable indexicals), both are anchored by a doxastic center, and both are *de se*. But the potential anchors for the attested shiftable indexicals must belong to one of a conventionally given class: e.g. Amharic *-ññ* can only be anchored to ©*, ©[@] or ©^{say}, the last center corresponding to the agent of a verb of saying. But EMAS may shift anchoring more promiscuously. They are more like *come* (Barlew 2015), in that they may be anchored to any contextually salient and RELEVANT perspective. So though both are indexical in my technical sense—anchored to a discourse perspective—in this respect, EMAS are more like anaphoric elements than like English indexicals as classically conceived (Kaplan 1977/1989).

Finally, note that the conclusions I draw in §4.5 about the influence in discourse of the assertion of EMA statements, and about the corresponding patterns of apt response, bear more generally on the role of evidentiality in discourse and on whether evidential statements across languages are modal or, e.g., involve an evidential speech act operator. The conclusions here suggest that evidentials may always involve epistemic modality as part of their meaning, a meaning that might be more specific for particular evidential particles in a variety of ways—e.g., specifying that the evidence the anchor brings to bear is of a certain type (direct, reportative, inferential, etc.). But it may be that in languages with pervasive evidential morphology, the evidentiality *per se* is never itself at-issue, even though the evidential's prejacent is itself not asserted (as in the EMA statements considered above). According to Greg Kierstead (p.c.), this is the case with Tagalog reportatives evidential *daw*. Are these cases where the evidential morpheme is proffered but *conventionally* backgrounded? What would that mean?

There is much more to say, but space precludes further speculation. I hope these notes will prove useful to others considering the rich class of epistemic modal auxiliaries, and epistemic modality and evidentiality more generally.

References:

- Abusch, Dorit (2002) Lexical Alternatives as a source of Pragmatic Presuppositions. In B. Jackson (ed) *Proceedings of SALT 22*. Cornell University Department of Linguistics.
- Abusch, Dorit (2009) Presupposition Triggering from Alternatives. *Journal of Semantics* 27: 37-80.
- Amaral, Patricia Matos, E. Allyn Smith & Craige Roberts (2008) Review of Potts: *Conventional Implicature*. *Linguistics and Philosophy* 30:707-749.

- AnderBois, Scott, Adrian Brasoveanu & Robert Henderson (2015) At-issue proposals and appositive impositions in discourse. *Journal of Semantics* 32 (1):93–138. doi: 10.1093/jos/fft014.
- Asher, Nicholas & Alex Lascarides (2003) *Logics of Conversation*. Cambridge University Press.
- Barlew, Jefferson (2014) Coming toward a doxastic agent: A doxastic analysis of the motion verb *come*. *Proceedings of Sinn und Bedeutung 19*. Georg August University at Göttingen. Göttingen, Germany.
- Barlew, Jefferson (2015) Perspectives in discourse. Ms., The Ohio State University.
- Beaver, David (2010) Have you Noticed that your Belly Button Lint Colour is Related to the Colour of your Clothing? In R. Bauerle, U. Reyle, and T. E. Zimmermann (eds.), *Presuppositions and Discourse: Essays offered to Hans Kamp*. Oxford: Elsevier, 65–99.
- Beaver, D. & Clark, B. (2008) *Sense and Sensitivity: How Focus Determines Meaning*. Oxford: Wiley-Blackwell.
- Brennan, Virginia M. (1993) *Root and Epistemic Modal Auxiliary Verbs*. Ph.D. dissertation, University of Massachusetts at Amherst.
- Cummins, C., Amaral, P., & Katsos, N. (2012) Experimental investigations of the typology of presupposition triggers. *Humana Mente* 23:1-16.
- DeRose, Keith (1991) Epistemic possibilities. *The Philosophical Review* 100(4):581-605.
- Janice Dowell (2011) A Flexibly Contextualist Account of Epistemic Modals. *Philosopher's Imprint* 2011, pp.1-25. https://www.academia.edu/1125714/A_Flexible_Contextualist_Account_of_Epistemic_Modals
- Ebert, Christian, Cornelia Ebert & Stefan Hinterwimmer (2014) A unified analysis of conditionals as topics. *Linguistics and Philosophy* 37(5):353-408.
- Eckardt, Regine (2014) *The Semantics of Free Indirect Discourse*. Brill, Leiden, The Netherlands.
- Egan, Andy (2007) Epistemic Modals, Relativism, and Assertion. *Philosophical Studies* 133:1–22.
- Egan, Andy, John Hawthorne & Brian Weatherson (2005) Epistemic modals in context. In G. Preyer & G. Peter (eds.) *Contextualism in Philosophy: Knowledge, Meaning and Truth*. Oxford University Press, 131-170.
- Faller, Martina (2002) *Semantics and pragmatics of evidentials in Cuzco Quechua*. PhD thesis, Stanford University.
- Farkas, Donka & Kim B. Bruce (2010) On reacting to assertions and polar questions. *Journal of Semantics* 27.
- von Fintel, Kai & Anthony S. Gillies (2007a) *Might* made right. In Andy Egan & Brian Weatherson (eds) *Epistemic modality*. Oxford University Press.
- von Fintel, Kai and Anthony S. Gillies (2007b) An opinionated guide to epistemic modality. In T.S. Gendler & J. Hawthorne (eds.) *Oxford studies in epistemology: volume 2*. Oxford: Oxford University Press, 32–62.
- von Fintel, Kai & Antony S. Gillies (2008) CIA leaks. *The Philosophical Review* 117 (1): 77–98.
- von Fintel, Kai & Anthony S. Gillies (2010) *Must*. . .stay. . .strong! *Natural Language Semantics* 18:351-383.
- von Fintel, Kai & Sabine Iatridou (2003) Epistemic containment. *Linguistic Inquiry* 34(2):173-98.
- von Fintel, Kai & Sabine Iatridou (2008) How to say *ought* in foreign: The composition of weak necessity modals. In J. Guéron & J. Lecarme (eds.) *Time and Modality*. Spring, 115-141.

- Ginzburg, Jonathan (2012) *The Interactive Stance: Meaning for Conversation*. Clarendon Press, Oxford.
- Glanzberg, Michael (2007) Context, Content, and Relativism. *Philosophical Studies* 136: 1–29.
- Glass, Leila (2013) Deriving indirectness and questioning entailment for epistemic *must*. Handout of a talk at the LSA, Boston, January 2013.
- Grice, H.P. (1957) Meaning. *The Philosophical Review* 66 (1957).
- Hacking, Ian (1967) Possibility. *Philosophical Review* 76(2):143–68.
- Hacquard, Valentine (2006) *Aspects of Modality*. Ph.D. dissertation. MIT.
- Hacquard, Valentine (2013) On the grammatical category of modality. In M. Aloni, M. Franke & F. Roelofsen (eds.) *Proceedings of the 19th Amsterdam Colloquium*.
- Harris, Jesse A. & Christopher Potts (2009) Perspective-shifting with appositives and expressive. *Linguistics and Philosophy* 32(6):523-552.
- Heim, Irene 1983. On the projection problem for presuppositions. In M. Barlow, D. Flickinger and M. Wescoat (eds.) *Proceedings of WCCFL 2*, Stanford University, 114-125.
- Kamp, Hans & Uwe Reyle (1993) *From Discourse to Logic*. Kluwer, Dordrecht.
- Kaplan, David (1977) Demonstratives, draft #2. Ms., UCLA Philosophy Department. Revised and published as "Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals," in Joseph Almog, John Perry & Howard Wettstein (eds.) *Themes from Kaplan*, Oxford University Press, 1989, pp.481-563.
- Karttunen, Lauri (1971) Some observations on factivity. *Papers in Linguistics* Volume 4, 55-69.
- Karttunen, Lauri (1973) Presuppositions of Compound Sentences. *Linguistic Inquiry* 4:169-193.
- Kierstead, Greg & Craige Roberts (2014) Evidentiality At Issue. Talk at the Questions In Discourse workshop, Stuttgart, Germany, May, 2014.
- Koev, Todor (2013) *Apposition and the Structure of Discourse*. Ph.D. dissertation, Rutgers, The State University of New Jersey – New Brunswick.
- Kratzer, Angelika (1981) The notional category of modality. In H. J. Eikmeyer and H. Rieser (eds) *Words, Worlds and Contexts*. de Gruyter, Berlin, pp.38-74.
- Kratzer, Angelika (1991) Modality. In A. von Stechow and D. Wunderlich (eds) *Semantics: An international handbook of contemporary research*. de Gruyter, Berlin, pp.639-650.
- Kratzer, Angelika (1998) Scope or pseudoscope? Are there wide-scope indefinites? In S. Rothstein (ed.) *Events and Grammar*. Kluwer, Dordrecht, 163-196.
- Kratzer, Angelika (2009) Building a pronoun: Fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry* 40:187–237.
- Lasersohn, Peter (2005) Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy* 28:643-86.
- Lasersohn, Peter (2008) Quantification and perspective in Relativist semantics. *Philosophical Perspectives* 22(1):305-337.
- Lewis, David (1979b) Attitudes *de dicto* and *de se*. *The Philosophical Review* 88.4:513-543.
- Lycan, William (2001) *Real conditionals*. Oxford: Oxford University Press..
- Lyons, John (1977) *Semantics*. Cambridge: Cambridge University Press.
- MacFarlane, John (2005) The assessment sensitivity of knowledge attributions. In T. Szabó Gendler & J. Hawthorne (eds.) *Oxford Studies in Epistemology I*. Oxford University Press, pp.197-233.
- MacFarlane, John (2011) Epistemic modals are assessment-sensitive. In B. Weatherson & A. Egan (eds.) *Epistemic Modality*. Oxford University Press, 144-178.

- Martin, Scott (2013) *The Dynamics of Sense and Implicature*. Ph.D. dissertation, The Ohio State University.
- Martin, Scott (2014) Supplemental update. Ms., Nuance Communications. Submitted for review.
- Moltmann, Friederike (2012) Two kinds of first-person-oriented content. *Synthese* 184:157-177.
- Moore, G. E. (1993) Moore's Paradox. In Baldwin, Thomas. *G. E. Moore: Selected Writings*. London: Routledge. pp. 207–212. ISBN 0-415-09853-X.
- Morgan, Jerry (1970) On the criterion of identity for Noun Phrase deletion. *CLS* 6:380-381.
- Moss, Sarah (2015) On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics* 8.5:1-81.
- Ninan, Dilip (2010) De Se Attitudes: Ascriptions and Communication. *Philosophical Compass* 5:551-67.
- Palmer, F. R. (1990) *Modality and the English Modals*, 2nd edition. Longman.
- Papafragou, Anna (2006) Epistemic modality and truth conditions. *Lingua* 116:1688-1701.
- Portner, Paul (2007) Beyond the common ground: The semantics and pragmatics of epistemic modals. In J.-Y. Yoon & K.-A. Kim (eds) *The Perspectives of Linguistics in the 21st Century*. Seoul: Hankook Publishing Company.
- Portner, Paul (2009) *Modality*. Oxford Surveys in Semantics and Pragmatics. Oxford University Press.
- Potts, Chris (2005) *The Logic of Conventional Implicatures*. Oxford University Press, New York.
- Ramchand, Gillian (2014) Stativity and present tense epistemics. *Proceedings of SALT* 24:102-121.
- Roberts, Craige (2007) Specificity. Handout, talk at UMass, Amherst.
- Roberts, Craige (2011) *only*: A case study in projective meaning. In Barbara H. Partee, Michael Glanzberg & Jurgis Skilters (ed.) *Formal Semantics and Pragmatics: Discourse, Context and Models*. Special issue of the *Baltic International Yearbook of Cognition, Logic and Communication*, Riga, Latvia.
- Roberts, Craige (2012b) Information Structure: Afterword to Roberts (1996). *Semantics and Pragmatics* 5.7:1-19. <http://dx.doi.org/10.3765/sp.5.7>.
- Roberts, Craige (2015) *Indexicality: de se semantics and pragmatics*. Ms., OSU.
- Roberts, Craige (2015b) Speech acts in discourse context. Ms., OSU.
- Roberts, Craige (in press) Accommodation in a Language Game. In Barry Loewer & Jonathan Schaffer (eds.) *A Companion to David Lewis*. Wiley, Hoboken, NJ.
- Roberts, Craige, Mandy Simons, David Beaver & Judith Tonhauser (2009) Presupposition, implicature and beyond: A unified account of projection. N. Klinedinst & D. Rothschild (eds.) *Proceedings of the ESSLLI 2009 Workshop on New Directions in the Theory of Presupposition*, Bordeaux.
- Saebø, Kjell Johan (2009) Judgment Ascriptions. *Linguistics and Philosophy* 32: 327–52.
- Schaffer, Jonathan (2011) Perspective in Taste Predicates and Epistemic Modals. In Andy Egan & Brian Weatherson (eds.) *Epistemic Modality*. Oxford: Oxford University Press, 179-226.
- Schlenker, Philippe (2003) A plea for monsters. *Linguistics and Philosophy* 26:29-120.
- Simons, Mandy (2006) Observations on embedding verbs, evidentiality, and presupposition. *Lingua* doi:10.1016/j.lingua.2006.05.006.
- Simons, Mandy, Judith Tonhauser, David Beaver & Craige Roberts (2010) What projects and why. In Nan Li & David Lutz (eds.) (2010) *Semantics and linguistic theory (SALT)* 20, 309–327. Ithaca, NY: CLC Publications.

- Simons, Mandy, David Beaver, Erich Kummerfeld, Craige Roberts & Judith Tonhauser. The best question: Explaining the projection behavior of factives. Accepted for a special issue of *Discourse Processes on Questions In Discourse*.
- Slote, M. (1978) Time in counterfactuals. *Philosophical Review* 87:3–27.
- Snyder, Eric (2013) Binding, genericity, and predicates of personal taste. *Inquiry* 56:278-306.
- Stalnaker, Robert (1974) Pragmatic Presuppositions. In M. Munitz and Peter Unger (eds.), *Semantics and Philosophy*. NY: New York University Press, 197-214.
- Stalnaker, Robert C. (2008) *Our Knowledge of the Internal World*. Oxford University Press. Chapter 3: “Locating Ourselves in the World”.
- Stalnaker, Robert (2014) *Context*. Oxford University Press.
- Stephenson, Tamina (2007) A Parallel Account of Epistemic Modals and Predicates of Personal Taste’. In Puig-Waldmüller, E. (ed.) *Proceedings of Sinn und Bedeutung 11*. Barcelona: Universitat Pompeu Fabra, pp. 583–97.
- Stone, Matthew (1994) The reference argument of epistemic *must*. *Proceedings of the International Workshop on Computational Semantics (IWCS)* 1:181-190.
- Swanson, Eric (2006) *Interactions with Context*. Ph.D. dissertation, MIT.
- Swanson, Eric (2011) How not to theorize about the language of subjective uncertainty. In Andy Egan & Brian Weatherson (eds.) *Epistemic Modality*. 2011, OUP, 249-269.
- Syrett, Kristen & Todor Koev (to appear) Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics*.
- Tonhauser, Judith (2012) Diagnosing (not-)at-issue content. *Proceedings of SULA 6*, UMass, Amherst.
- Tonhauser, Judith, David Beaver, Craige Roberts & Mandy Simons (2012) Towards a taxonomy of projective content. *Language* 89.1:66-109.
- Judith Tonhauser, David Beaver, Judith Degen, Marie-Catherine de Marneffe, Craige Roberts & Mandy Simons (2015) Negated evaluative adjective sentences: What projects, and why? Talk to be given at Experimental Pragmatics 2015, University of Chicago, July 16-18.
- Veltman, Frank (1985) *Logics for conditionals*. PhD thesis, University of Amsterdam.
- Wechsler, Stephen (2010) What ‘you’ and ‘I’ mean to each other: Person indexicals, self-ascription, and theory of mind. *Language* 86.2:332-365.
- Yalcin, Seth (2007) Epistemic modals. *Mind* 116:983-1026.
- Xue, J., & Onea, E. (2011) Correlation between presupposition projection and at-issueness: An empirical study. In *Proceedings of the ESSLLI 2011 Workshop on Projective Meaning*. Ljubljana, Slovenia.