Craige Roberts OSU & NYU

Agreeing and Assessing: Epistemic Modals and the Question Under Discussion

Oslo Workshop on "Non-At-Issue Meaning and Information Structure "
June 30, 2017

Abstract: Important debates in the recent literature on Epistemic Modal Auxiliaries (EMAs) hinge on how we understand disagreements about the truth of assertions containing EMAs, and on a variety of attested response patterns to such assertions. Some relevant examples display evidence of faultless disagreement ((Lasersohn 2005; Egan et al. 2005; MacFarlane 2005, 2011; Egan 2007; Stephenson 2007) or "faulty agreement" (Moltmann 2012). Others display a variety of patterns of felicitous response to statements with EMAs, responses which sometimes seem to target the prejacent alone, and other times the entire modal claim (Lyons 1977; Swanson 2006; Stephenson 2007; von Fintel & Gillies 2007b,2008; Portner 2009; Dowell 2011; among others). I provide an alternative characterization of what it is to agree about EMA statements. Then I provide evidence that the pattern of felicitous response to 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 the course of this exploration, I motivate a revision of the notion of RELEVANCE (from Roberts 1996/2012), which makes it sensitive not only to the simple truth or falsity of a target proposition, but to its possibility or probability as an answer to the QUD.

1. Introduction[1]

I'm going to focus on two kinds of issues recently considered in the literature on the semantics of epistemic modal auxiliaries: issues about the meaning of an assertion whose main clause includes an epistemic modal auxiliary (**EMA**) like *might* or *must* as reflected in responses to such assertions. The relevant data have been used to call into question the usual formal semantic treatment of EMAs in natural language since Montague, wherein they are treated as modal operators that take a propositional complement (the **prejacent**) to yield a proposition that is true

[1] Acknowledgements: This paper has developed slowly, over several years. In developing it I owe a debt to several colleagues with whom I've worked on closely related projects. Foremost, I am grateful to Judith Tonhauser, Mandy Simons and David Beaver for many hours of work on what it means to be at-issue and on the identification of the Question Under Discussion; and to Jefferson Barlew, Greg Kierstead and Eric Snyder, with whom I worked on a project on the role of perspective in interpretation. The work of the first group has benefitted from support from the US National Science Foundation grant #1452674, "What's the question? A cross-linguistic investigation into the question under discussion". The second group received support from a Targeted Investment in Excellent grant and a Research Enhancement Grant from the Colleges of the Arts and Sciences at The Ohio State University, and from grant #0952571 from the National Science Foundation, "Modeling Perspectival Expressions". This work also benefitted enormously from a Fellowship in 2014-15 at the Institute for Advanced Studies of Central European University in Budapest, and especially from extended discussions there with Zsofí Zvolensky. I presented related work with Greg Kierstead at the Questions In Discourse workshop in 2014 in Stuttgart, Germany; and earlier versions of the present paper were presented at the Research Institute for Linguistics of the Hungarian Academy of Sciences, Budapest and at the MASZAT working group there; at ReDraw 2015 at the University of Groningen, Netherlands; at a workshop on Differentiating Contexts at Carnegie Mellon University in December, 2015; at a workshop on Questions in Pragmatics at the ILLC at the University of Amsterdam, 2016; and in 2017 at the CUNY GradCenter Cognitive Science Colloquium; at the UMass/Amherst Linguistics Department colloquium; the Yale University Linguistics and Philosophy Colloquium; the UIO-XPrag.de Workshop on Non-At-Issue Meaning and Information Structure in Oslo; and a colloquium at the Universitat Pompeu Fabra, Barcelona. I also thank two anonymous referees and Michael Franke at *Semantics and Pragmatics*. I am very, very grateful for all this support and feedback. Of course, all errors are my own.

just in case the prejacent is true in some (for *may*, *might*, etc.) or all (for *must*, *should*, etc.) of the worlds in a contextually restricted set of possible worlds.

The first of these issues is about what it means to agree or disagree with someone who has made an assertion with an utterance whose main clause includes an epistemic modal auxiliary. This is illustrated by an example due to Stalnaker (2014: Chapter 6):

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

Stalnaker uses the example to argue against a proposal due to Yalcin (2007), where (roughly) under attitude predicates an index over information states shifts to the set of possible worlds compatible with the agent's attitude. On Yalcin's account, the complement of *believe* in (1), *it might rain*, is interpreted relative to Jones' beliefs: so here, the complement means that in Jones' belief state there is at least one world in which it rains. But Stalnaker points out that on such an interpretation, we are hard-pressed to explain what Alice is disagreeing with. Presumably she's not taking issue with her own claim that it follows from Jones' beliefs that it may rain. Related examples have been used by Relativists like Egan (2007) and MacFarlane (2005, 2011) to argue that the proposition expressed by an EMA is not itself anchored to the agent's (here Jones') belief state—that is to say, it isn't Jones' beliefs that provide evidence of the truth of the modal proposition. Instead, it expresses a simple possibility, but is interpreted with respect to a distinguished parameter of interpretation, the *judge* or *assessor*. On that story, Jones' belief and Alice's involve a disagreement about the same proposition—that there's the simple possibility of rain.

This is intuitively plausible from the point of view of explaining (dis)agreement, and especially its "faultlessness": Even if Jones has incorrect information, his belief state should, indeed, properly entail the possibility of rain. But there are a number of problems with the Relativist approach to such examples. For one thing, variants on (1) with quantified agents are acceptable:

(2) Alice: Everyone else in my group believes that it might rain, but I disagree—I think we can count on fine weather all day.

The problem is how to permit an index of interpretation to vary over judges that are quantified over in the object language. Lasersohn (2008) offers an operator which does just that, but Snyder (2013) argues that his account is both empirically inadequate—over-generating and predicting that examples may have strange truth conditions—and conceptually *ad hoc*; and Kneer (nd.a) offers additional arguments against this way of resolving the issue, as well as experimental evidence that many of the readings predicted by Relativist accounts are not generally attested by native speaker informants. And there are other problems with Relativism.[2]

[2] For example, for embedded contexts like that in (1), the Relativist approach requires special context shifters ("Monsters"). Though those have been argued for recently for the semantics of so-called "shifted indexicals", e.g. in Schlenker (2003), others have argued that that approach is empirically inadequate (e.g., Anand & Nevins 2004, Deal 2013) for the indexicals. And the introduction of Monsters constitutes a pretty big departure from standard assumptions about context, arguably under-motivated by the issue of faultless disagreement.

But if we take the Relativist approach to be inadequate, the problem for the standard semantics remains: What does Alice disagree with, if it is not the truth of the proposition expressed by *it might rain* in that complement? If we can explain the intuitions about disagreement for free, from the pragmatics, the resulting theory would be simpler and less *ad hoc*.

The second issue is illustrated by von Fintel & Gillies' (2008:83-84) Mastermind examples. These consist of (constructed) dialogues between players of the game Mastermind. The game involves a pair of players, one of whom, the codemaker, has full access to an array of colored pegs; the other player tries to guess the colors and pattern within a limited period. Play 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:

(3) [Pascal and Mordecai are playing Mastermind. After some rounds where Mordecai gives Pascal hints about the solution, Pascal says:] There might be two reds.

(4) [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 (4a) and (4c) clearly respond to the evidential claim in (3)—they affirm or deny whether Pascal correctly reported his own evidential state, while (4b) and (4d) respond to the prejacent alone. If Mordecai is playing his codemaker role properly, we expect a response that gives Pascal feedback about his own evidence (a or c) instead of announcing the correct solution (b or d). For example, response (4c) 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 truth of the simple prejacent alone. This clearly involves a solipsistic anchoring of *can't* to Pascal, the same anchoring as that of *might* in (3); we understand that Mordecai himself, knowing the correct solution, might not be able to truthfully assert (3), nor does his response in (4a) or (4c) pertain to whether *he* takes it to be possible that there are two reds: he might know that in fact there is only one red. Rather, he is responding to what is possible *given what (he knows that) Pascal knows at that point in the game*.

However, more to the point here, though the responses in (4b) or (4d) would spoil the game, they are still perfectly felicitous responses to Pascal's (3), or to a question like *Might there be two reds?*. And in those cases, these assessments would be understood to target the truth of the prejacent alone, 'there are two reds' (or the polar question 'are there two reds?'), not the entire EMA (or corresponding modal question). The problem devolves from the possibility of assessing the truth of the prejacent alone.

Lyons (1977) argued that in cases where the prejacent alone is being assessed, the modal acts as a speech act operator (cf. Faller's 2002, 2006 treatment of Quechua evidentials). Swanson (2006) claims that in such cases the modal merely serves to performatively raise the possibility that the

prejacent is true (while in others it has something more like its standard semantic effect); and von Fintel & Gillies (2007a,b) speculate that such EMAs might involve two speech acts. But as von Fintel & Gillies (2008) point out, the speaker who proffers the original EMA in such a case can still reasonably defend himself against a claim of having said something false: *I didn't say there ARE two reds, I said there might be so far as I knew.* Any approach to the puzzle which discounts the truth conditional contribution of the modal incorrectly predicts that the speaker has said something false. Thus, whatever the merits of a speech act approach to evidentiality generally, it fails in the EMA cases to offer an adequate explanation of the attested response patterns and intuitions about speaker commitment.

Accordingly, many assume that evaluative responses like those in (4) must target the full compositional meaning of the EMA. Hence, they assume that the attested responses tell us something direct about the truth conditional content of the EMAs being responded to, and in particular about the semantic contribution of the epistemic modal auxiliaries.

But both the kinds of issues for EMAs addressed here provide evidence that interlocutors, in considering the meaning of an utterance, are not so constrained: they needn't always assess the full proposition expressed. In contexts where the prejacent is one of the alternatives under consideration at the time an EMA is issued, both agreement and assessment take into consideration the broader pragmatic consequences of the utterance for the evolving context, including the way in which the prejacent would be reflected in that context were the assertion accepted. That is, in responding to an EMA, interlocutors are negotiating update of the common ground and other features of the context of utterance, especially including implications of the utterance for the Question Under Discussion (see below).

My argument, then, will be that though there are problems with the standard semantics of epistemic modality, the phenomena associated with agreement and response do *not* pose problems for that semantics, and instead should be explained pragmatically, with a view to what a speaker using or responding to an EMA aims to accomplish: comparing or calibrating belief states, and contributing to the strategy of inquiry in which an EMA occurs.

This essay is organized as follows: In §2, I offer a brief overview of my assumptions about the role of context in interpretation, including the Question Under Discussion (QUD). §3 offers a brief introduction to the way that anaphoric evaluative responses like *yes* and *no* are understood in discourse, and discusses an empirical generalization over their relationship to the provocations they respond to, which I call Apt Response. §4 discusses what it means to agree with an EMA, while §5 discusses what it means to assess such an assertion relative to the QUD. §6 offers a possible revision of the notion of RELEVANCE to the QUD. And §7 offers conclusions and prospects for future work.

2. Background: Context and interpretation

The notions and theoretical assumptions outlined in this section have been independently motivated in the course of addressing issues orthogonal to those that are our main interest in this paper. Thus, the assumptions on which this account will depend are non-ad hoc.

2.1 Context, the QUD and RELEVANCE

First, a terminological note: In what follows I'll often talk about the **anchor** of an EMA, as I have informally above. I'm using the term to pick out the agent whose judgment is reflected in the EMA. In a Relativist account, this would be the Judge with respect to which the proposition denoted is interpreted. In a Contextualist account, this will play a regular part in the proposition expressed; for example, under a Kratzerian approach (Kratzer 1981), it would be the agent whose belief state or evidence serves the modal base for interpretation of the prejacent of the EMA. I intend the term to be used in a relatively theory-neutral way here. It is quite clear that though the speaker is often the anchor of an EMA (Hacking 1967, DeRose 1991, Egan et al. 2005, MacFarlane 2006), it may also be the addressee (Fintel & Gillies 2007a, especially in "interrogative flip"), the agent of an embedding attitude (Stephenson 2007) or some other party (von Fintel & Iatridou 2003, von Fintel & Gillies 2007a, Dowell 2011)—especially in Free Indirect Discourse (Roberts 2015a); see Roberts (2015a) for extended discussion of how to capture the available anchors in a given context.

I adopt the following notion of the context of utterance, simplified from Roberts (1996/2012, 2004, 2015b):

(5) The context of utterance in a discourse D, $<DG_D, QUD_D, CS_D>$, consists of
- DG: the interlocutors' evident goals, their Domain Goals
- QUD: the ordered set of questions currently under discussion, with the first ("current question") CQ. These constitute the interlocutors' discourse goals: addressing these questions.
- CS: the interlocutors' Context Set, the set of worlds compatible with their common ground (in the sense of Stalnaker 1978).

As usual in contemporary semantics, I take a proposition to be a set of possible worlds, so that the CS is the intersection of the set of propositions in the CG. A question is a set of sets of worlds—intuitively the set of possible answers to the question, formally a partition over the CS. An answer to a question is an assertion that entails the falsehood of at least one cell in the partition, one answer; a complete answer eliminates all but one cell.
`

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. Here is how that was characterized in Roberts (1996/2004/2012):

(6) **RELEVANCE**: 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 CQ.
- b. A question is RELEVANT to a QUD iff it has an answer which contextually entails a partial or complete answer to the CQ.
- c. A directive is RELEVANT to a QUD iff its realization promises to play a role in resolving the CQ.

Hence, a RELEVANT utterance **addresses** the QUD, entails that at least one possible answer is wrong, bringing us closer to resolving it.

At any given point in discourse, the current question under discussion may be part of a larger strategy of inquiry, addressing some larger question. A strategy of inquiry is a constrained organization of questions designed to break down an over-arching QUD into more manageable sub-questions. Respecting RELEVANCE, sub-questions and the questions they subserve are ordered under entailment, in the sense of erotetic entailment defined by Groenendijk & Stokhof (1984): q entails q' just in case every complete answer to q entails a partial answer to q'. See Roberts (1996) for formal definitions and illustrations.

2.2 Proffered content and at-issueness[3]

(7) The semantic **CHARACTER** of a word or expression consists of several types of content, including:
Presupposed content: that which constrains the contexts of utterance in which utterance of the content is felicitous
Auxiliary content:[4] that which is directly attributed to the speaker and/or some other agent whose point of view is contextually salient
Proffered content:[5] the compositionally calculated truth conditional content of the expression; what it contributes to what is asserted, asked or directed by an utterance in which it occurs

The semantic contribution of a constituent, its CHARACTER, consists of all the constituent's semantic contents, including what's presupposed, its proffered content, and any auxiliary content.

For the purposes of discussion, I make the standard assumption that a **proposition** is a set of possible worlds. A **question** is a partition on the CS, the cells in the partition (each cell a proposition) constituting the set of possible answers to the question.

Now we can define a notion that will play an important role in what follows:

[3] This material section was developed in close collaboration with David Beaver, Mandy Simons, and Judith Tonhauser.
[4] Auxiliary content includes the content associated with Potts' (2005) Conventional Implicatures. Amaral et al. (2007) and Harris & Potts (2009) provide evidence that such content is anchored to a salient point of view, frequently, though not necessarily, that of the speaker or addressee. See recent treatments in Anderbois et al. (2015), Martin (2014), and Barlew (2015).
[5] This is intended to correspond to Heim's (1982,1983) asserted content, generalized as in Roberts (1996) to cover the non-presupposed content of non-declaratives; and to Rooth's (1996) Ordinary Semantic content.

(8) **at-issueness**:[6]
If a proposition *p* is part of the proffered content of an utterance *U* and *p* is RELEVANT to the QUD at the time of *U*, then *p* is **at-issue** in *U*. Otherwise, it is **not at-issue**.

This definition is revised from that in Simons et al. (2011) with the addition of the term *proffered*, thus explicitly excluding auxiliary content and content which is merely presupposed from what can potentially be at-issue.[7] Thus, non-proffered content cannot by itself satisfy the requirement that an utterance be RELEVANT to the QUD. This is consistent with Potts' (2005) claim that appositives and other C(onventional)I(mplicature)s, which contribute auxiliary content, are "nondeniable". Amaral et al. (2007) extend this observation to demonstrate that we normally cannot directly respond to an appositive like that in (9) with an anaphoric affirmation or denial.

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

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

In (9), 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 the less direct *Hey! Wait a minute!*, as in (10). Note that *Hey! Wait a minute!* is not itself evaluative; it is merely a conversation stopper, and the subsequent interjection must be explicit if it's to target the appositive.

The same kind of evidence argues that some proffered content is not-at-issue in certain contexts, as my colleagues and I have argued.[8] For example, consider what happens when proffered content is backgrounded via the prosodic Focus of the utterance in which it occurs:

Focal backgrounding:
(11) [A, B, and C are discussing the race that B and C just watched, in which their mutual friend Mary was supposed to be a participant.]
A: How did Mary run?
B: She ran [QUICKLY]$_F$—one of her best times ever!
C: No she didn't!
C′: Hey, wait a minute! Mary didn't run at all! / But Mary didn't run at all!

[6] This version of at-issueness was developed in close collaboration with David Beaver, Mandy Simons, and Judith Tonhauser, as part of our long-term project to investigate the character and role of the QUD in interpretation. To appreciate its development and general role in interpretation, see Simons et al. (2011), Tonhauser et al. (2012), Simons et al. (2016) and Beaver et al. (2017).
[7] I have argued elsewhere (Roberts 2011) that some content, like the prejacent of *only*, can be both proffered and presupposed; I think that is the case with factive complements, as well. Such content is potentially at-issue, by virtue of being proffered.
[8] See the work cited in the previous footnote.

A's question presupposes that Mary ran and asks how. Then B's answer can be understood to have two parts:[9] the **theme** is that part that corresponds to the background assumption of the question posed: *Mary ran*; while the **rheme** is the "new" content, the part that gives the answer: *quickly*. As usual (Rooth 1992, Roberts 1996, Schwarzschild 1999), the rheme is prosodically focused, bearing accent, while the theme is prosodically backgrounded, not bearing accent. In this context, C's answer can only be understood to bear on the rhematic portion of B's response, denying that the way Mary ran was quickly. It cannot be taken to deny that Mary ran. To deny that thematic portion of the utterance, C would have to say something like one of the utterances in C′.

3. Provocation and evaluative response

(9) – (11) illustrate an empirical generalization I will call Apt Response; I mean here *apt* in the sense of 'having a tendency', 'ordinarily disposed' (Merriam-Webster https://www.merriam-webster.com/dictionary/apt). This generalization bears on how sequences of a **provocation** and an evaluative anaphoric **response** in discourse are pragmatically constrained.

(12) A provocation/response sequence is a sequence of speech acts (each a question, assertion, or direction) such that the second responds to the first.

To show how this works, we'll first briefly consider the class of anaphoric evaluative responses, including the English polarity particles *yes* and *no* and their kin.

(13) An **anaphoric evaluative response**
- (a) is anaphoric to the provocation, presupposed to be a speech act that can be evaluated as true or false, i.e. a polar question or assertion, and
- (b) evaluates the provocation by assessing its truth.

The provocation in such a sequence is a speech act (question, assertion, direction), and the anaphoric evaluative response it prompts is truth-evaluative, like those in (9)-(11) above.

Then we'll characterize Apt Response, and finally show how it and RELEVANCE (as defined in the previous section) together constrain response interpretation.

3.1 Evaluative responses and polarity particles

The class of anaphoric evaluative responses includes *yes*, *no*, *maybe*; *uh-huh*, *huh-uh*; *yup*, *nope*; *that's true*, *that's false*; *that's right*, *that's wrong*; *that's (in)correct*, etc.[10] The polarity particles *yes/no* and their less formal cousins *uh-huh/huh-uh* and *yup/nope* may respond to a polar question or an assertion (via addressing the corresponding yes/no question). The remaining

[9] See Roberts (2011) for discussion of the theme/rheme terminology, which I use informally to refer to parts of the semantic content of an utterance that play different roles in its information structure.

[10] I take it that the provoking question or assertion is the antecedent for the anaphoric response. It is argued by some that *yes* and *no* are elliptical, rather than merely anaphoric (Kramer & Rawlins 2011), but since I take it that ellipsis is a special type of anaphora (Roberts 2016), that difference is not relevant for present purposes.

responses, *that's true/false*, *that's right/wrong* and *that's (in)correct* usually respond to an assertion. In other words, all these respond to a speech act which can be evaluated as true or false

Note that *ok* is also anaphoric but is not an evaluative response in this sense: *ok* conveys agreement with a salient proposal. When used to respond to a speech act provocation, it can signal agreeing to accept the commitment to action proposed by a directive (in my terms, agreeing to add it to the target agent's set of evident intentions—see Roberts 2015b), agreement to add an assertion to the CG, or agreement to address a question—to add it to QUD. While agreeing to add the proposition expressed by an assertion to the CG *amounts to* accepting its truth, that relationship to truth is indirect. It is the use of *ok* to respond to directions, not possible for the truth-evaluative anaphoric responses above, that tells the difference.

All the above differ from simple propositional anaphors, such as the propositional use of *this* or *that*. The evaluative response anaphors presuppose as antecedent not a simple proposition, but a discourse proposal: an assertion or question which can be evaluated as true or false. Similarly, *ok* in the sense just discussed presupposes a discourse proposal, which may also be a directive.

Farkas & Bruce (2010) offer an account of the role of English polarity particles *yes* and *no* in discourse. The empirical facts are summarized in my Table 1 (though *provocation* and *response* are not their terminology); each cell in the last two columns is illustrated by the examples in (15) below, which all involve a simple provocation with a subject and intransitive predicate. Each of the responses in (15) has a short form—the polarity particle alone, and a full form, followed by a clause clarifying the proposition which the respondent intends thereby to commit to. For convenience, I call such a clarificatory clause the particle's **coda**. The provocation/response combinations in the last two columns of Table 1 may felicitously include either the short or the full form of the response noted. By "intended coda" I mean to suggest that if the response is in short form, it is to be understood as if it had a coda with the indicated polarity.

(14) Farkas & Bruce's (2010) features for a provocation/response sequence involving a polarity particle:

[+] :	property of main clause of provocation or coda: positive polarity
[-] :	property of main clause of provocation or coda: negative polarity
[SAME] :	property of the provocation/response sequence: polarity of the intended coda (*he is/isn't*) matches that of the denotation of the sentence radical of the provocation (matrix sentence of an assertion, sentential complement of an interrogative)
[REVERSE] :	property of the provocation/response sequence: polarity of the intended coda (*he is/isn't*) doesn't match that of the denotation of the sentence radical of the provocation

response polarity: provocation polarity:	particle polarity	intended coda [+]	intended coda [-]
provocation [+]	absolute	[SAME] *yes*: (15a+i), (15b+i)	[REVERSE] *no*: (15a+iii), (15b+iii)
provocation [-]	absolute	[REVERSE] *yes*: (15c+i), (15d+i)	[SAME] *no*: (15c+iii), (15d+iii)
	relative	[REVERSE] *no*: (15c+iv)	[SAME] *yes*: (15c+ii)

Table 1: Confirming ([SAME]) and Reversing ([REVERSE]) Moves in Provocation/Response Pairs with Polarity Particles in the Response
after Farkas & Bruce (2010)

(15) Potential provocation/response pairs:
Provocation:
Positive polarity: (a) Sam is home.
(b) Is Sam home?
Negative polarity: (c) Sam isn't home.
(d) Is Sam not home?
Response:
Positive particle: (i) Yes / Yes, he is.
(ii) Yes / Yes, he isn't.
Negative particle: (iii) No / No, he isn't.
(iv) No / No, he is.

As you see, in all the examples involving absolute polarity, the particle takes as its antecedent the proposition expressed by the entire assertive provocation or provoking polar question, and either confirms it/answers positively or denies it/answers negatively. Absolute polarity displays the full paradigm, each provocation potentially responded to with a bare polarity particle with the sense of either (i) or (iii).

Farkas & Bruce say that relative polarity "signals sameness", by which I take it they mean that the polarity of the particle signals confirmation [SAME] if it accords with the polarity of the provocation ("the denotation of the sentence radical of the provocation") or reversal [REVERSE] if it does not, whatever the provocation's polarity. At first glance, the Table might suggest that relative polarity is more constrained than absolute, only licensed by provocations with negative polarity. As reflected there, Farkas & Bruce only give examples with a negative indicative provocation. However, note that with positive provocations the absolute and relative interpretations collapse: to confirm a positive assertion, we use *yes*; to reverse it, we use *no*. But there is no such collapse in response to a negative assertion. Relative polarity occurs readily in responding to a negative assertion, and in those cases may felicitously involve only the bare polarity particle without the coda. So (15c+ii), with bare *yes* in the response (and appropriate prosody, etc.), is a felicitous confirmation, despite the fact that the polarity of the provocation fails to match that of the understood coda. And (15c+iv) with bare *no* as response, may be

understood as a denial,[11] despite the fact that the particle's polarity is the same as that of the (negative) provocation.

But when the provocation is a negative question, typically (e.g., in the absence of some other gesture or facial expression to indicate the polarity mismatch) a response where the polarity particle is intended to have a relative interpretation must include an overt coda.[12] So the bare polarity particles (in (i)) are infelicitous, whereas the full forms, with coda, are felicitous:

(16) a) for [- + reverse] (15d+iv): Is Sam not home?
i) #No
ii) ✓No, he is.
b) for [- - same] (15d+ii): Is Sam not home?
i) #Yes
ii) ?Yes, he isn't.
iii) ✓Yeah no. ['he isn't'][13]

Sadock & Zwicky (1985:190) explain the necessity of a coda in responding to related examples like *Isn't Sam home?* in terms of the well-known bias of English questions with wide-scope negation, signaled in this variant by the initial contracted *isn't*:

> [T]he English question *Isn't it raining?* can be used not just to ask whether it is not raining, but simultaneously to indicate that the speaker guesses that it is. A simple positive answer could be very confusing. It could either be interpreted as a positive response to the question itself ('Yes, it is not raining') or as agreement with the speaker's guess ('Yes, you're right; it is raining').

I think this is right, and it suggests an interesting possibility. Suppose that negative questions, even those without overt indication of wide scope, generally suggest that the speaker is inclined to accept the positive answer to the polar question (or, perhaps in some cases, might have expected it to be the answer)—this implication constituting "the speaker's guess". Then in the relative polarity interpretation of *yes* or *no* in response to a negative question, the speaker agrees or disagrees, respectively, with the implicated positive proposition, as indicated by the coda. That is, what they confirm or reverse is the proposition that the provocation implicates to be more likely to be true. Then in these uses, the polarity particles act like what Sadock & Zwicky (1985) call an *agreement system*, like that of Japanese, indicating agreement or disagreement with what is said, or in this case implicated.

The understanding of relative polarity just proposed drives a wedge between assessment and agreement: when the provocation is a negative assertion, the respondent indicates that s/he agrees (*yes*) or disagrees (*no*) with the original speaker about the proposition expressed by the clause under the scope of negation, and in this case that is expressed in a way that's inconsistent with

[11] Farkas & Bruce adopt a definition of *denial* after van der Sandt & Maier (2003), according to which a full denial is a move that commits the author to the complement of the denotation of a provoking declarative sentence. Denials may also be partial, denying only part of what has been asserted.
[12] These are my judgments, these combinations not discussed in detail by Farkas & Bruce, except where noted.
[13] This is Farkas & Bruce's (50), with the interpretation I believe they intend suggested by the implicit coda.

the meaning of an absolute polar response. We might say that **the absolute use of polarity particles is a use for assessment**—of the truth of the proposition expressed by the provocation, whereas **the relative use is a use for agreement**.

3.2 Apt Response

But things are somewhat more complicated than the picture of anaphoric response just sketched would suggest. As we saw in (9) – (11), a provocation may involve more than one proposition. When we respond with a bare anaphoric response, with no clarifying coda, we depend on prior discourse to make it clear which proposition is being responded to.

(17) **Apt Response**: An evaluative anaphoric response will be understood to target the at-issue content of the provocation. I.e.
- **Apt Response to Question**: if the provocation is a question with at-issue content *q*, the response addresses *q*.
- **Apt Response to Assertion**: if the provocation is an assertion with at-issue content *p*, the response addresses *?p*.

Apt Response is a strong empirical generalization, reflected in felicitous patterns of provocation/response that we'll consider shortly. This tendency is even stronger when the evaluation is **purely anaphoric**, i.e. as in (9) or (10) where it occurs without a clarifying coda, or when a coda doesn't disambiguate between the proffered and non-proffered content of the provocation, as in (11). These reflect most clearly the role of salience in anaphora resolution. One hypothesis about anaphora resolution which has independent support from the experimental work in psycholinguistics (see extended discussion in Roberts 2016) is the following:[14]

(18) **at-issue (proffered+RELEVANT) content is more salient than not at-issue content**. Hence, at-issue content is more likely to be the target of anaphora. (Roberts 2016)

Attending to the QUD helps focus attention. When interlocutors can predict what each other will be attending to at any given point, they can predict what they'll all take to be salient. This is especially important for retrieving meaning when anaphora of any kind is involved, including in anaphoric responses of the sort under consideration here. If I just say *no* or *yes*, what is the default assumption about which polar question or assertion I'm responding to in all of preceding discourse? Even assuming that this was a response to a particular provocation, that provocation itself may have been complex, as with the modal/prejacent cases. In some cases where there are two options for some reason, using a coda can disambiguate; but when we use the simple anaphoric response, interlocutors just have to figure out which would have been intended in the context of utterance. Which? And how do we know?

Why would we expect at-issue (and hence RELEVANT) content to be more salient than not at-issue content? I think the answer lies in this direction (which Grosz (1977) was hinting at long ago): A speaker tacitly intends such anaphoric elements to resolve to that content which

[14] Note that not all at-issue content is equally salient. E.g. the prejacent of *only* is lexically backgrounded relative to its exclusive implication, and therefore tends to project. When it *is* at-issue, it doesn't project, but that does not mean that it is as readily available for evaluative anaphoric response as the exclusive implication (Roberts 2011).

interlocutors (i) could reasonably think that she was attending to and (ii) could take her to reasonably have expected them to be attending to as well. Insofar as she can reasonably predict what her interlocutors would assume in this regard, she can reasonably expect those interlocutors to retrieve her intended anaphoric antecedent. Then we can explain central patterns for resolution in the empirical data (as reflected in both the experimental literature and corpus studies), on the assumption that competent interlocutors (a) are attending to what's RELEVANT in the proffered semantic content of the utterance—i.e. to its at-issue content—and can expect each other to be doing so, as well, and (b) hence can felicitously respond anaphorically to that at-issue content and expect other speakers to do so as well, all other things being equal.

Accordingly, I think the following is a natural extension of the role of RELEVANCE as defined above:

(19) **extended RELEVANCE**: Since what is crucial to the felicity of a provocation is its RELEVANCE TO the (preceding) QUD, the default assumption about what's important in evaluating the provocation is assessing whether its at-issue (hence RELEVANT) content is true.

By assessing the at-issue content of a provocation, the interlocutors determine whether they take it to help resolve the QUD, furthering their discourse goals.

3.3 The role of RELEVANCE and Apt Response in response interpretation

Apt Response and the requirement of RELEVANCE yield a set of predictions about how a sequence of provocation and evaluative response will be interpreted, summarized in Table 2:

Character of Interaction		Predictions as a function of	
Provocation	**Response**	**the content of the Provocation**	**the content of the Response**
Question	Answer	**Apt Response to Question:** only the at-issue content of a Question can be addressed by a bare anaphoric answer ex.(20)	**Answer RELEVANCE:** only the proffered content of an Answer can satisfy RELEVANCE to a Question ex.(21),(22)
Assertion	Evaluation	**Apt Response to Assertion:** only the at-issue content of an Assertion can be targeted by a bare anaphoric evaluation ex.(9)-(11)	**Evaluation RELEVANCE:** only the proffered content of an Evaluation plays a role in evaluating the Assertion ex.(23)
Constraint follows from:		Apt Response	RELEVANCE

Table 2: **Predictions about Provocation/Response interpretation**

In the third column, Apt Response to Question and Apt Response to Assertion pertain to the range of felicitous evaluative anaphoric responses to a provocation in a given context, and predict which aspect(s) of the content of the provocation an evaluative anaphoric response can felicitously target. When the provocation is a question, Apt Response predicts that only the

proffered-and-RELEVANT content of the question plays a role in how we understand an anaphoric answer. This is illustrated in (20), wherein the yes/no question itself has both non-proffered and proffered content:

(20) 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 appositive *a fearless leader* in (20A) is auxiliary content (Potts 2005 and much subsequent work), and hence cannot be at-issue. Thus, we understand *yes* or *no* in B's response to target at most the proffered content; hence the infelicity of tags that target the non-proffered content in (20B′). The question is whether Edna has started the descent; and what counts as a felicitous direct response differs between B and B′ accordingly.

Apt Response to Assertion is illustrated by (9)-(11) above. When the provocation is an assertion, its at-issue content will be more salient than its not at-issue content, hence a purely anaphoric evaluative response will only be understood to target the at-issue content. So only the proffered, potentially at-issue content in A's utterance in (9A)/(10A) can be the target of the purely anaphoric denial in (9B); to target the auxiliary content, something like the non-direct *Hey! Wait a minute!* response is required.[15] In (11), though proffered, the presupposed and prosodically backgrounded proposition that Mary ran is not at-issue, and hence not taken to be the target of the direct anaphoric response *no*; note that *she didn't* fails to disambiguate between 'she didn't run quickly' and 'she didn't run'.

Among the RELEVANCE-based patterns in the fourth column of Table 2, Answer Relevance is illustrated for CI content by (21).[16] A *wh*-question cannot be felicitously addressed by an utterance in which the content of an auxiliary CI is crucial to ensuring the response's RELEVANCE to the question:

(21) Where's Edward?
a. #Edward, who's in Santa Cruz, loves to surf.
b. Edward, who loves to surf, is in Santa Cruz.

In (21a), the main clause without the appositive, *Edward loves to surf*, does not by itself address the QUD posed by the explicit question *Where's Edward?*. Since only proffered content can satisfy the requirement of RELEVANCE, and the non-restrictive relative clause, though RELEVANT, is not proffered, (21a) is infelicitous. In (21b), the contents of subject and relative clause are switched, yielding a RELEVANT at-issue main clause, *Edward is in Santa Cruz*, and the result is felicitous.

In addition, the focally backgrounded portion of an answer cannot be understood to be the rhematic portion of a response to a question. (22) is the Answer Relevance counterpart of (11):

[15] See Tonhauser's (2012) carefully constructed diagnostics (1a, 1b, 1c), designed for using Apt Response to Assertion to explore what's at-issue in fieldwork elicitation or through controlled experiment.
[16] See Tonhauser's (2012) Diagnostic 2.

(22) Who's Edward Witten?
#[EDWARD WITTEN]$_F$ is now the top-dog in string theory.
Edward Witten [is now the TOP-DOG IN STRING THEORY]$_F$

Finally, Evaluation Relevance is illustrated for CIs by (23), where (23B) is a very odd response to (23A) because the only content in (B) that's relevant to (A), justifying (B)'s rejection, is the auxiliary content of a CI:

(23) 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 reflecting these constraints argue that appositives are not part of the proffered content of an utterance in which they occur. Hence, they are not potentially part of the at-issue content of the utterance.[17] But the constraints apply not only to the auxiliary content of CIs, but to other types of not at-issue content as well, including content which is proffered—contributing to the truth conditional content of the clause in which it occurs—but prosodically backgrounded, as we saw in (11). Something similar can be shown for other types of conventionally backgrounded content, like the non-clefted portion of a cleft sentence.

Apt Response also reveals a parallel between response patterns and patterns of presupposition and projection, especially interesting in cases where triggers fail to consistently lead to projection. For example, factive verbs like *know*, *discover*, and *be aware of* are said to presuppose the truth of their complements. 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 sometimes factive complements fail to project as predicted; Abusch (2002, 2009) thus dubs them "soft presupposition triggers". Beaver (2010) offers numerous naturally occurring examples where 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,[18] whether with respect to an explicit interrogative or, as in the following naturally occurring example, due to an implicit QUD:

(24) …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 (24), 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

[17] Under a different definition of what it is to be AI, there are at-issue uses of non-restrictive relative clauses. See Syrett & Koev (2015). But those occur under special circumstances and will not concern us here.
[18] As noted above, this might lead one to argue that the complement of a factive is proffered; whether or not it is also conventionally presupposed is an open question.

environment. See Beaver (2010) for many other examples, Simons et al. (2015) for further discussion, especially of cases like (25) where the failure of a factive complement to project in the first sentence in (25B) is a function of its rhematic status, as reflected in surface prosodic FOCUS:

(25) A: James just found out that Harry's having a graduation party, and I just can't understand why he's so upset about it.
B: He didn't find out that [HARRY]$_F$'S having a graduation party, …
he found out that [HARRIET]$_F$ is having a graduation party, and [HARRIET]$_F$ is his best friend.

See also 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 to the QUD, 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 (non-rhetorical) QUD is felicitous, the CG does not entail its complete answer (Roberts 1996). Therefore, in such contexts the complement does not project, a behavior only licensed by pragmatic presupposition. There is now experimental evidence that at-issueness is key to understanding projection with factives (Xue & Onea 2011; Tonhauser, Beaver & Degen ms.) and with the normally factive evaluative adjectives as well (Tonhauser et al. 2015).

But the other side of this coin, not often discussed, is that when the complement of a factive is at-issue, the question of whether the attitude associated with the factive obtains between the agent and the complement is typically *not* itself at-issue. This 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 as parentheticals. 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:

(26) 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 (26a-c) are evidential. With some embedding verbs, this type of use licenses the slifting we see in (26b), wherein the subject and verb are postposed, rather like an adverbial. In (26a-c), the proposition denoted by the

complement, 'Louise has left town' is directly RELEVANT to the explicit QUD. This is reflected in the possibility of an anaphoric evaluative response to the truth of that complement, (26d).[19]

Note that even though the attitude itself is parenthetical in these uses, that does not mean that the attitude predicate plays no role in the truth conditions of the proposition expressed: the speaker in uttering one of (26a) – (26c) is not thereby committed to the truth of the complement. Thus, the attitude in one of these assertions is still proffered, and acts like an evidential hedge. Denying the evidential statement instead of the complement itself, as in (26e), weakens the proffered evidence for the truth of the RELEVANT complement, and is thereby indirectly RELEVANT itself. In this respect, proffered but not at-issue parenthetical predicates are quite different from the CIs we considered earlier, since the latter cannot felicitously license direct response, as we saw in (9). But in (26d) and (26e), there is a helpful coda clarifying which of the propositions in the proffered content of (26a-c) is being targeted—the entire proposition (26e) or the (non-entailed) complement (26d).

4. Agreement

4.1 Epistemic answers

Epistemic modal answers may have a special status with respect to the QUD: clarifying the inquiry by making evident one possible answer—the modal's prejacent. For example, consider the following dialogue:

(27) Ames: [clueless new detective, just joining the homicide team, knows nothing about the investigation under way. Speaking to the lead investigator Brand:]
I hear you've made some progress in the case. Who's the murderer?
Brand: It might be the neighbor.

Semantically a question like Ames's is modeled as a partition over the interlocutors' Context Set (CS): the set of answers to the question which are compatible with the Common Ground. But

[19] 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 do not do so, even in non-parenthetical, embedded uses. 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 (24) above, where the semi-factive *discover* occurs in the antecedent of a conditional. 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 (iA), where a range of parenthetical semi-factives could be felicitously used:

(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, (24) and (iA), share is that the complement is at-issue relative to the explicit QUD. So, as with the non-factive parentheticals *believe* and *think* in (26), we can also directly respond to (iA) with (26d). Thus, the thesis that at-issue content fails to project, evident in (24), converges here with evidence from Apt Response and RELEVANCE that the complements of parenthetical semi-factives are at-issue.

technically, Brand's reply is not RELEVANT to Ames's question, in that it fails to contextually entail the falsehood of any of the cells in that partition. Why then is Brand's response intuitively relevant to the QUD posed by Ames?—as it clearly is.

Notice that in the set-up for (27), we specified that Ames is relatively ignorant about the investigation. In particular, he presumably didn't know in advance of asking the question who the possible suspects were. Brand's response serves to clarify the question, asserting that one of the cells in the partition is the prejacent, as well as thereby implicating that in other cells the prejacent is false—if there were only one live-option, then the question would be answerable, and presumably the cooperative Brand would so indicate.

We often use EMAs in this way—to clarify the question, making evident the alternatives which remain open in an inquiry. This is related to Swanson's (2006.§2.3:73ff) claim that on some occasions someone saying *might p*, serves to performatively raise the possibility that *p*. Swanson claims that such EMAs can have two other functions: asserting the whole modal proposition, and advising the addressee(s) about how to change their subjective probabilities for the prejacent of the modal. This seems intuitively correct from a functional point of view. But it 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 (as Swanson assumes)?, and (b) Are there systematic contextual constraints on which function we can take *might p* to serve in a particular utterance? If there are such systematic contextual constraints, 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 in reply to question (a).

It's clear from context in some examples, like (28), that the doxastic (belief) state associated with the EMA is part of what's addressed in the response:

(28) [Watson to Holmes, both privy to the same body of evidence:] Then the butler must be the murderer!
[Holmes 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. This disagreement is about what the interlocutors' joint evidence entails, and accordingly about how their belief states should agree. This is functionally important in the inquiry they're conducting, because it pertains to what possible answers to the main question (*Who's the murderer?*) remain open.

Functionally, this is just like the standard MasterMind responses by Mordecai, in (4a,c). Such cases, then, clarify the nature of the QUD at a certain point in an inquiry.

Then what is it to agree or disagree in such a case? I think it is exactly the same thing that's at issue in Alice's disagreement with Jones in (1), above. Consider also what Moltmann (2012) calls "faulty agreement", illustrated by the following:

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

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

In (29) 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 (30). But on Contextualist accounts, 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?

One more type of example helps to clarify how agreement with an EMA is not about the truth value of the proposition expressed. Suppose the young detective Ames from (27) above looks over the evidence and comes up with his own hypothesis about what happened:[20]

(31) Ames: Why are we focused on the neighbor? I think the wife might be the murderer.
Brand: I disagree. It's highly unlikely, given her alibi.

Note that Ames' initial question suggests that he takes the investigation's focus on the neighbor to be inappropriate, thereby implicating that the suspect he then suggests, the wife, may be a more likely suspect. Note further that the *it* in Brand's response seems to refer to the prejacent of Ames' EMA, 'the wife is the murderer'. We don't take Brand to be saying that the *possibility* of the wife being the murderer is unlikely—he could even explicitly affirm that possibility without contradiction: *it's possible, but highly unlikely, given her alibi*. Thus, Brand's response does not target the whole EMA. It is evaluative, but not truth-evaluative: the disagreement does not lie in an assessment of the *truth* of the EMA (with or without solipsistic anchoring to the speaker) *or* of its simple, non-modal prejacent, but in a judgment of the prejacent's implied likelihood.[21] We might say that Ames' EMA has raised the question of whether the prejacent is true, 'is the wife the murderer?'. That seems to be Ames' intention—to discuss whether that polar question should be overlooked as a possible answer to the over-arching QUD in the inquiry, 'who was the murderer?' and even to implicate that it is at least as likely as the hypothesis that the neighbor is the murderer. Then Brand, too, turns to discuss that question; his response suggests that pursuing that question probably wouldn't be the most fruitful sub-inquiry in the investigation. In other words, what Brand is disagreeing with isn't the truth or possibility of the prejacent, but its implied likelihood.[22]

[20] My understanding of this example benefited from discussion with participants at my presentation in the CUNY CogSci Speakers Series in October, 2016, and especially the probing questions of Dan Harris and Una Stojnic.
[21] See Moss (2015) on the semantics of graded epistemic modals like *likely*.
[22] Thanks to Zoltán Szabo (p.c.) for his sharp questions about the role of implication in this example.

Patterns of felicitous response also argue that agreement cannot be reduced to a negative assessment of the truth conditional content of a provocation. Sundell (2011:281) argues "From the fact that speakers disagree—whatever we take disagreement to be—it simply does not follow that the speakers literally express inconsistent propositions." So consider (32) and (33), each with alternate responses B and B′ to the provocation A:

(32) A: I really love the flavor of this hundred year old egg!
B: I disagree. I think it's terrible.
B′: #That's false—I think it's terrible.
B″: #Nuh-uh, I hate it!

(33) A: Given the information I received from the Sargent, the wife might be the murderer.
B: I disagree. She's got a good alibi.
B′: #That's false. She's got a good alibi.
B″: #Nuh-uh—she's got a good alibi.

In both (32) and (33), B′ can only be understood to deny the truth of the proposition expressed by A. This is not what B means—the responder is not arguing that A doesn't love the egg, or that A doesn't think the prejacent of A is possible. In neither case can the short anaphoric portion of the response in B′ convey what B does; I mark these as infelicitous because the follow-up clause then seems irrelevant to the denial. The attitude with which B disagrees in (32A) is non-doxastic—love is not about belief but about something like dispositions and desires; while the EMA in (33A) reflects a doxastic attitude. But the distinction here between the attitudes of A and B seems fully parallel, as does the inappropriateness of the reply in B′. The answers in B″, inspired by Sundell's observations about *nuh-uh*, are similarly infelicitous—the negative polarity of the particle assesses the proposition expressed by A; but then the follow-up is irrelevant: the issue is not what A's attitude toward the egg's flavor should be (32) or about which hypotheses A's information at the time of her utterance was consistent with.

(1), Moltmann's paradigm in (29) and (30), (31), and (32)/(33) together strongly argue that (dis)agreement isn't essentially about the truth of the proposition expressed by the speaker in a speech act or even the truth of a modal prejacent or the complement of an attitude predicate, but about a feature of the belief states of the relevant parties. What the premises in (29) and (30) and the complement in (1) 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 all those agents 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. In (1) Alice disagrees with Jones in that while his belief state admits of that possibility, hers does not. And in (31), the disagreement pertains to the likelihood of the truth of the proposition 'the wife is the murderer'. **This is what agreement and disagreement are about: belief states that are or are not congruent in the relevant way**.

4.2 Agreement and attitudes

To make this a bit clearer, we define:

(34) A **likelihood ranking** for agent *x* over a set of propositions P is a partial order $<_{x,P}$ such that for all p, p′ ∈ P, p $<_{x,P}$ p′ just in case *x* takes it to be more likely that *p* is true than that *p*′ is true.[23]

(35) The **belief state** of agent *x* in world *w* at time *t*, DOX(x,w,t) is an ordered pair $<W_{DOX(x,w,t)}, <_{x,P}>$ such that:
- $W_{DOX(x,w,t)}$ is the set of possible worlds in which all the propositions are true that *x* in *w* at *t* takes to be true, and
- $<_{x,P}$ is a likelihood ranking over P = {p | ∃w∈$W_{DOX(x,w,t)}$: w∈p}, i.e. the set of all propositions *x* takes to be possibly true.

The factors bearing on the likelihood ranking $<_{x,P}$ would presumably include, e.g., whether the proposition in question is true in all worlds in $W_{DOX(x)}$ or only some (and if some, how many/what proportion), the strength of the evidence making various propositions seem possibly true, etc.[24]

Given this rough characterization, then, we can define a notion of doxastic congruence:

(36) Belief states b and b′ are **congruent** with respect to proposition *p* (*p*-congruent) just in case the two states concur with respect to the truth, possibility, necessity and likelihood of *p*, i.e.:
p is true in b [true in all worlds in b] iff it's true in b′
p is possible in b [there is some world *w* in b s.t. *p* is true in *w*] iff it's possible in b′
p is evidently necessary in b [true in all worlds *w* in b in which the agent's possibly indirect evidence is true][25] iff it's evidently necessary in b′
p is (un)likely in b to the nearly the same degree that it's (un)likely in b′, as reflected in the likelihood rankings in b and b′

(37) Belief states b and b′ are ***p*-possibility-congruent** just in case they concur with respect to the possibility of *p*. Similarly for ***p*-evidently-necessary-congruent**, ***p*-truth-congruent**, and ***p*-likelihood-congruent**.

(38) Agents **agree** with respect to a belief about proposition *p* to the extent that their belief states are congruent with respect to *p* in the relevant respects.

[23] Instead of basing the relative likelihood of a given pair of propositions on this primitive ranking, one could instead base it on a comparison of the subjective probability distribution (in a given context) over those propositions, e.g. as defined in Davis, Potts & Speas (2007).
[24] Of course, this characterization of an agent's belief state in a given world and time is inadequate and implausible. For one thing, there is the problem of logical omniscience (as noted in footnote 15 above): we do not even recognize all the propositions that follow from, or are consistent with, what we believe (or believe to be possible), let alone rank them all for relative likelihood. What matters here is that part of an agent's belief state is a partial ranking over the evident and relevant possibilities.
[25] Inspired by von Fintel & Gillies' (2010) evidential characterization of *must*, wherein an EMA containing it needn't carry the same degree of epistemic certainty associated with simple assertion of the prejacent.

Agreement so-characterized cannot be reduced in the general case to shared assessments of the simple truth of a proposition—'what is said'. Belief states as characterized above are richer than propositions, and they admit of a richer notion of agreement and disagreement: about possibility, necessity and likelihood, as well as simple truth or falsity. So in (1), Alice disagrees with Jones about the possibility of rain. But in (28) Holmes disagrees with Watson about the necessity of the butler being the murderer—though they presumably agree about its possibility. In (31), Brand agrees with Ames about the possibility that the wife is the culprit, just not about its implied likelihood.

In fact, disagreement in the most general sense is richer still, going beyond the doxastic, as has been discussed in the relevant philosophical literature. López de Sa (2015:158) argues that "[T]he ordinary, intuitive notion of disagreement is flexible and covers cases that may well go beyond that of *contradictory* judgments with respect to one and the same content and judgments with *contradictory* contents. . ." MacFarland (2014) talks about "practical disagreement", illustrated with the following example:

(39) Suppose that Jane likes Bob, but Sarah hates him. In a perfectly respectable sense, Jane disagrees with Sarah, even if she believes all the same things about Bob. She does not disagree with Sarah about whether p, for any *p*, but she disagrees with Sarah about Bob, since Sarah's attitude towards Bob is not cotenable with hers. (MacFarland 2014:122)

The notion of *non-cotenability* describes both practical disagreement and doxastic disagreement—the disagreement over the truth of some *p*. Two attitudes *a* and *a′* are non-cotenable if an agent who holds attitude *a* could not coherently adopt *a′* without dropping *a*. Another case of non-cotenable attitudes is due to Huvenes (2011):[26]

(40) Suppose that Pierre and Marie are conducting a series of experiments in order to test a scientific hypothesis that they both believe to be false and that the following conversation takes place after one of the experiments:

Pierre: *The hypothesis is false.*

Marie: *I disagree, we need to do further testing.*

In this case we do not need to understand Marie as denying the content of the sentence uttered by Pierre. Indeed, since she herself believes that the hypothesis is false, and thus accepts the content of the sentence uttered by Pierre, it is more plausible to take her to disagree with Pierre on the grounds that there is insufficient evidence for him to assert that the hypothesis is false without further testing. In other words, Pierre and Marie disagree even though Marie is not denying the content of the sentence uttered by Pierre. (Huvenes 2011:170)

We might describe the disagreement here as a difference in level of confidence about the falsity of the hypothesis in question—we could say that Pierre and Marie assign different degrees of subjective probability to the falsity of the hypothesis. Note that Marie cannot express her disagreement with Pierre by saying either *yes* or *no, that's true* or *that's false*. Similarly, Dreier

[26] Huvenes tells us that Weatherson (2009) also notes this sort of disagreement, in debates about non-cognitivism in moral philosophy.

(2009) characterizes disagreement in language in a way that comes close to the present proposal, taking it to be based on a more fundamental disagreement between the mental states expressed.[27]

Many others in the philosophical literature on disagreement cite the multiple ways in which two agents may disagree, including Baker (2014), Egan (2014), Huvenes (2012), Kölbel (2014), Marques (2014), Stojanovic (2007) and Sundell (2011). Nonetheless, most of the authors in this literature persist in what Plunkett & Sundell (2013:246) call "a crucial and largely unargued for premise about disagreement":[28]

> The premise is that the best way to explain why an exchange between two speakers expresses a genuine disagreement is, in almost all circumstances, to suppose that those speakers mean the same thing—that is, express the same concepts—with the words they use in that exchange.

For example, Egan (2014:90-92) distinguishes between "disagreement in thought" and "disagreement in language" and says that disagreement in discourse involves "trying to get our interlocutors to update their beliefs in incompatible ways". But in the end, he characterizes disagreement in language as "two interlocutors asserting something that's incompatible with each others' beliefs", so that their "conversational moves are in conflict". So ultimately he aims to explain the disagreement relevant in the EMA statements under consideration here in terms of their truth conditional content.[29]

As philosophers often note, disagreement can be both a state and an activity. The phenomenon of interest here is not fundamentally the activity, wherein two or more agents engage in a dispute over the correct view on some matter, but the state, wherein two or more agents fail to concur about some attitude, and in particular when they hold noncotenable attitudes. This is not, of course, to say that the two kinds of disagreement are unrelated. When two individuals dispute the truth of some proposition, then if sincere, this is evidence that they hold non-cotenable doxastic attitudes, and hence are in a state of disagreement. Thus, if agent A issues a denial of some proposition explicitly expressed by agent B, A purports to be in a state of disagreement with B: thus, agreement and assessment tend to collapse in absolute evaluative anaphoric responses to simple, non-negative assertions or polar questions. But as we saw, in many more cases

[27] Marques (2014) also talks about disagreement in terms of "attitude incompatibility", but then focuses on the semantics of the propositions expressed.

[28] Though space precludes a thorough discussion of this literature, let me note that the authors do adopt a variety of approaches to the problem of how to reconcile disagreement and faultlessness, which go well beyond relativism. Some of these authors restrict their study to examples involving explicit linguistic disagreement; for example, Stephenson (2007) limits herself to disagreement reflected in the use of *no* and *nuh-uh*, though she admits that there may be a broader phenomenon. And not all reduce the problem to the proffered content of the utterances in question. For example, López de Sa argues for a presuppositional approach, and Sundell (2011) ultimately adopts something like a metalinguistic approach inspired by Barker (2002)—involving disagreement over standards for the application of predicates of personal taste. Nonetheless, even they do take a linguistic approach to understanding disagreement, whereas here I am arguing that it is fundamentally non-linguistic, albeit often linguistically expressed.

[29] Egan takes the propositions expressed by the relevant examples to be essentially *de se*, and he concludes (p.94) that "Difference in *de se* belief—even possession of incompatible *de se* beliefs—isn't in general happily described as disagreement." In other words, he specifically rejects as unintuitive a (centered) version of the type of proposal I offer here, but without empirical argument. The intuition isn't clear to me, and I would say it is inconsistent with the linguistic evidence.

disagreement and denial—negative assessment—come apart, as in the difference between an absolute vs. a relative evaluative anaphoric response to an assertion with negative polarity, and as attested by the infelicity of negative polarity responses (a form of assessment) to express disagreement in examples like (29), (30), (31) and (40).

Hence, I take the recent debate over faultless disagreement prompted by utterances involving EMAs seems generally to be somewhat misguided, since the principals tend to assume that what agents disagree about is truth simpliciter—the truth of the EMA provocation, or some other aspect of the meaning of expressions uttered. Similarly, one can argue that retractions and disavowals are basically about one's former beliefs, rather than about the truth of the statements that reported them. One might speculate that this is more generally the case, pertaining to the debate over predicates of personal taste as well—these are about congruence of taste, not judgments of the truth of statements about taste. See Kneer's (nd.b) experimental evidence about predicates of personal taste.

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 prejacents (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 saw, we do just that in cases involving group investigations based on joint evidence—as in (28) pertaining to Holmes' and Watson's case—and in games like Mastermind—as in Mordecai's responses (4a, c). But if that is so, how do we know, in a particular case, just what the disagreement is about?

(Dis)agreement in a particular discourse, and other types of judgment and response as well, are not necessarily directed at the entire proposition expressed by the proffered content of the target utterance, but at what's at-issue in that utterance. This is what's RELEVANT in the context of utterance and hence what interlocutors are most likely to be interested in reaching an agreement about, so that the QUD to which they're committed is satisfactorily addressed. Since QUD theory tells us that interlocutors carefully track what's at-issue at any point in a discourse, in a non-defective context this gives them adequate grounds for understanding the disagreements expressed. Then the response patterns we considered in §3.3 give us a way of testing this hypothesis. Felicitous evaluative responses 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.[30] Now we turn to consider what that might tell us about the status of the prejacent in EMAs vs. that of the full modal proposition itself.

5. Assessing EMAs with respect to the QUD

von Fintel & Gillies (2007b) notice that some uses of EMAs, as in their (41), display parallels with Simons' parenthetical uses of predicates, as in (26) above:[31]

[30] The structure of discourse is complex, and accordingly, responses can serve a variety of functions, not all of them evaluative (or directly so). Hence, the hypothesis developed here is *not* intended to serve as a complete theory of the semantics and pragmatics of response. See Asher & Lascarides (2003), Ginzburg (2012), and Farkas & Bruce (2010) for other important considerations.

[31] They do not analyze the example in terms of what's at-issue and RELEVANCE, but I take what they say to be compatible with the present explanation.

(41) QUD: Quinn: Why isn't Louise coming to our meetings these days?
a. Alex: She might/must be too busy with her dissertation.
Prejacent: 'Ann is too busy with her dissertation to come to our meetings'

Note that the QUD here presupposes that Louise isn't coming to the meetings in question. That goes unchallenged, so we can assume that all the worlds in the interlocutors' CS are worlds in which Louise isn't attending. Then the QUD partitions the CS to reflect the different possible reasons why she isn't attending. In this context, only the prejacent of Alex's response in (41a), understood as such a reason, is directly RELEVANT to the explicit question posed by Quinn, entailing the rejection of at least one possible answer. Hence, just as the complements are at-issue in the parentheticals in (26a-c), the prejacent is at-issue in (41a). The modal contributed by the EMA is pragmatically backgrounded, Alex's utterance serving as a means of drawing attention to the prejacent as one possible answer to Quinn's question. But just as (26a)-(26c) fail to entail their complements, the proffered content of Alex's response in (41a) does not entail the truth of the prejacent. Alex is not committed to the truth of the foreground, at-issue prejacent in (41a), evidence that the modal has its usual semantics.

As we saw above in (31), responding to a question with an EMA whose prejacent *p* is a possible answer to the question thereby both draws the addressee's attention to that possible answer—making it evident when it might not have been before—and in turn suggests that the interlocutors might profitably discuss the subquestion 'whether *p*', ***?p***. So Alex's reply in (41a) tacitly raises the sub-question 'has Louise not been coming because she's too busy with her dissertation?'. Roberts (1996) argues that questions are quite often implicitly raised, so long as the resulting strategy of inquiry is well-formed, as it is here. And much subsequent work supports that contention (Büring 2003, Simons et al. 2015, *inter alia*). In other words, (i) being too busy is a possible reason why someone doesn't attend a meeting, so (ii) the prejacent in Alex's response is at-issue, and the response (iii) tacitly raises the question whether the prejacent is true. Thus, Apt Response to Assertion leads us to correctly predict that responses to (41a) may target its prejacent *qua* answer to the QUD, as in (41b), a perfectly felicitous direct response:

(41) b. Rob: No she isn't—she filed it last month!
c. Rolf: That can't be right! I saw her out partying several times in the last month.

The structure of the interchange in (41a)+(41b):

not-coming: 'Louise isn't coming to our meetings'
busy: 'Louise is too busy with her dissertation to come to our meetings'
busy presupposes *writing*:
writing: 'Louise is writing her dissertation'

presupposed:	CS entails *not-coming*
Quinn's QUD:	*why not-coming?*
Alex's assertion:	$might^{\text{DOX(Alex)}}(busy)$
Implicated sub-question:	?*busy* with the resulting strategy of inquiry: <?*busy*, *why not-coming*>
Rob's assessment:	¬*busy*
Rob's evidence:	Louise has filed her dissertation

Note that Rob's Assessment is RELEVANT to the QUD, since it resolves the sub-question *?busy*, i.e. it tells us that one possible answer to the QUD, *busy*, is false.

Of course, (41c) is also felicitous. It's not *prima facie* obvious in this example whether *that* refers to the prejacent alone or the entire evidential claim conveyed by the EMA. But note that *might* in (41a) and *can't* in (c) are not plausibly anchored to the same agent. *might* is anchored to Alex. But we understand *can't* to be anchored to Rolf: Rolf isn't arguing that the prejacent *busy* is not consistent with (or doesn't follow from) Alex's belief state, but instead that it is incompatible with his own, an inconsistency for which he then provides evidence. That is, instead of directly assessing what Alex has said, in asserting (41c) Rolf disagrees with Alex: what can't be right is that the prejacent is the reason why Louise isn't coming to our meetings. Hence, the interlocutors are negotiating not about the truth of Alex's EMA statement *per se*, but about how to address the QUD and update the CG. Such disagreement entails an unwillingness to accept Alex's suggested explanation—the prejacent as a (possible or evident) answer to Quinn's question.

The structure of the interchange in (41a)+(41c):

not-coming: 'Louise isn't coming to our meetings'
busy: 'Louise is too busy with her dissertation to come to our meetings'
busy presupposes *writing*:
writing: 'Louise is writing her dissertation'

presupposed:	CS entails *not-coming*
Quinn's QUD:	*why not-coming?*
Alex's assertion:	$might^{\mathrm{Dox(Alex)}}(busy)$
Implicated sub-question:	?*busy* with the resulting strategy of inquiry: <?*busy*, *why not-coming*>
Rolf's disagreement:	$\neg[can^{\mathrm{Dox(Rolf)}}(busy)]$
Rolf's evidence:	Louise has been partying

Summarizing: By rejecting the *prejacent* with (41b), Rob negatively assesses that possible answer to the QUD, which amounts to answering *no* to the implicit ?*busy*. On the other hand, Rolf in (41c) disagrees with Alex's claim that *busy* is possible, his disagreement based on his own evidence. Again, this is RELEVANT to both questions in the strategy of inquiry, ?*busy* and the over-arching QUD, because asserting that the prejacent is incompatible with his evidence entails that so far as Rolf knows it's false, in turn entailing a negative answer to ?*busy*.

I note that (41a) involves an ellipsis, the tacit infinitival complement of *too*: 'too busy to come to the meeting'. But there is nothing essential about that ellipsis. We see the same pattern in (42) (after von Fintel & Gillies 2007a), without any elliptical content:

(42) QUD: 'what was the weather like yesterday when we were out of town?'
Alex: It might have rained.

Where, again, we find the felicitous responses:
Barbara: No, it didn't rain.

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

Examples like (41) and (42) are quite different from the Holmes/Watson example (28) and the Mastermind responses in (4).

(28) [Watson to Holmes:] Then the butler must be the murderer!
[Holmes to Watson, patiently:] No, my dear Watson. You're forgetting that the housekeeper also had the keys to the wine cellar.

Since Watson and Holmes are discussing their joint evidential state, the evidentiality and force of the EMA are RELEVANT, and Apt Response to Assertion 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*. Even if it happens that the butler *is* the murderer, this is not what's at-issue.

The structure of the interchange in (28):

butler: 'the butler is the murderer'

QUD:	*who is the murderer?*
Watson's assertion:	$must^{\text{Dox}(H\oplus W)}butler$
Holme's assessment:	$\neg(must^{\text{Dox}(H\oplus W)}\ butler)$

Similarly, following (3), what's at-issue in (4) is what Pascal's evidential state should be, not the truth of the prejacent *per se*, predicting under Apt Response to Assertion that directly responding to the prejacent as in (4c) doesn't commit Mordecai to any direct belief about the prejacent:

(3) [Pascal says:] There might be two reds.
(4c) [Mordecai replies:] That's wrong. There can't be.

The structure of the interchange (3)+(4c):

two-reds: 'there are two reds'

QUD:	*How many reds are there?*
Pascal's assertion:	$might^{\text{Dox(Pascal)}}$ *two-reds*
Mordecai's assessment:	$\neg(might^{\text{Dox(Pascal)}}$ *two-reds*$)$

Note that the possibility that the whole EMA is what's at-issue, just illustrated, differentiates English EMAs from evidentials in languages in which those are grammaticalized. In particular, in a language like Cheyenne, as discussed in detail in Murray (2010, 2014), evidentials themselves are never part of the at-issue content, and cannot license direct response. This seems to be typical of evidentials across a range of languages (e.g., see Faller 2002; Matthewson, Rullmann & Davis 2007; Lee 2011), which is presumably why they are sometimes treated as "speech act operators" (Faller 2002).

Now briefly consider how EMAs work in interrogative utterances, and the role they play under Apt Response to Question in determining what answers are felicitous and what those answers can be taken to mean:[32]

(43) [Context: Alex and Billy are roommates:]
Billy: Where are my keys?
Alex: I have no idea. Where might you have left them?
Alex's question: ?$where_x$[$might^{\text{DOX(Billy)}}$*Billy left keys at x*]

Billy's question in (43) (in the default case where the question is neither rhetorical nor a test-question) implicates that she doesn't know where her keys are. Then Alex's question in response displays 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. Hence we have the felicitous response in (43′):

(43′) Billy: I don't know. ¬*know*(*billy*,?$where_x$[$might^{\text{DOX(Billy)}}$*Billy left keys at x*])
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, and I haven't been anywhere else since I got home. (list of rejected possible answers)

But compare Alex's question in (43) and that in (44):

(44) Alex: Where are your keys?
Alex's question: ?$where_x$[*Billy's keys are at x*])
Billy: I don't know. ¬[*know*(*billy*,?$where_x$[*Billy's keys are at x*])]
But I'm pretty sure they're either in the car or they fell out of my purse in the restaurant. (complete list of possible answers)

Billy's two replies *I don't know* in both (43) and (44) are elliptical—meaning something like 'I don't know the answer to your question', so here anaphoric on the content of Alex's question. Thus Apt Response to Question predicts that as used in the two cases, responding to different questions posed by Alex, they do not denote the same proposition. In (44), *I don't know* means 'I don't know where my keys are', but the follow-up suggests that Billy does know where they might be; whereas in (43′), the reply means 'I don't know where I might have left them', presumably entailing that Billy doesn't know where the keys might be. Hence, the truth of *I don't know* in (44) is consistent with the falsity of *I don't know* in (43′).

However, anchoring an EMA to the addressee in an interrogative—the interrogative flip we just observed in Alex's response in (43)—is not automatic, but is another reflection of the context of utterance. This is illustrated in the following Mastermind case:

[32] (43) is inspired by an example in von Fintel & Gillies (2007a). Would that there were space to analyze their original example in detail here. See Dowell (2011) for enlightening discussion.

(45) [After some rounds where Mordecai gives Pascal hints about the solution, Pascal asks:] Might there be two reds?

Pascal's question: ?(*might*$^{\text{DOX(Pascal)}}$ *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 prejacent itself. So in (45) the only goal-RELEVANT anchor of *might* is Pascal, the speaker himself, resulting in a question with an EMA which does *not* involve interrogative flip. Then Mordecai can respond as follows:

(45′) [Mordecai:] No. ¬(*might*$^{\text{DOX(Pascal)}}$ *two-reds*)

Parallel with Mordecai's response as analyzed in (3)+(4c) just above, (45′) *no* means not 'no there aren't two reds' but 'no, you're wrong in concluding that there being two reds is consistent with the evidence you have now'. So RELEVANCE to the domain goal (wherein Pascal guesses the correct answer) gives us the anchoring for the EMA in (45), and Answer Relevance and Apt Response constrain interpretation of the response in (45′): the truth of the prejacent itself doesn't directly bear on the question.

The pattern of anaphoric evaluative responses to interrogatives that contain EMAs shows us that the flexibility of the anchoring of the EMA is independent of its foreground/background status—the anchor may be the addressee (flip) or the speaker. But in either case, asking an EMA question, as opposed to questioning the prejacent alone, is likely to make the modal content 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 the generalization Apt Response, and the resulting constraints on provocation/response relations, correctly predict attested patterns of response and their interpretations in discourses involving EMAs. Bear in mind that this account of the oft-attested difference between the modal-evidential implication and the prejacent in EMA assertions is non-ad-hoc, since the consequences of proffered content being at-issue relative to the QUD are attested in other types of content—CIs, prosodically backgrounded content, semi-factives and several other attitude predicates, the prejacent of *only* (Roberts 2011), etc. Even when content is not entailed by what a speaker says, as with the complement of parenthetical *believe* in (26) or the prejacents of the EMAs considered in this section, it may felicitously be the target of an anaphoric evaluative response so long as that content is what's at-issue and the entire utterance in which it occurs bears on the possibility that that content is a (possibly partial) answer to the QUD. Intuitively, being at-issue means being under discussion, and so, of course, interlocutors can respond to such content even if it hasn't been explicitly or directly asserted.

6. RELEVANCE Revised

We can understand why a speaker might cooperatively assert EMA-*p* even when only *p* itself, and not the evidential state reported, is an answer to the QUD. Observance of the maxim of Quality holds her back: 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, as in Simons' parenthetical attitude examples). Then what is the result of acceptance by the interlocutors of the solipsistically anchored assertion that EMA-*p*? First, as in Stalnaker (1978), 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, according to Roberts (1996), if *p* is RELEVANT, *?p* is RELEVANT, and *vice versa*.) *?p* is a partition over the CS, felicitous only if *?p* is RELEVANT to the previous QUD and the answer to *?p* is not entailed by the CG (so that both *p* and ¬*p* are live options).

Similarly, asking whether *might*-(*must*-) *p*—whether *p* is consistent with (follows from) the evidence available to the anchor—interrogatively implicates that so far as the speaker knows *p* is an open answer, and in turn suggests addressing the sub-question ?*p*.

This argues that when *p* is a potential answer to the QUD, assertion of EMA-*p* (or asking whether *p* is possible) is itself at least indirectly RELEVANT, because it suggests a strategy of inquiry for the QUD: addressing the sub-question ?*p*.

Thus, the kinds of examples we've considered here might be taken to motivate a reconsideration of what it is to be RELEVANT, so that we take into account not only the simple truth or falsity of possible answers to the QUD, but their possibility or likelihood, as well.

Consider (46):

(46) Q: Is it raining?
A: (i) It might be: People are coming in with wet umbrellas.
(ii) Well, John thinks so.
(iii) The weather report said there was a 2% chance of rain today.
(iv) They haven't cancelled the company picnic.

None of the answers in (46A) 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, via presenting evidence that bears on it. 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 (ii), the weather report is more often than not correct (iii), 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.

We could modify the definition of RELEVANCE from Roberts (1996) to guarantee that such utterances are directly RELEVANT. Here is a first stab at how we might do that (with help from Mandy Simons, p.c.):[33]

(47) An assertion that p is **RELEVANT** to question Q just in case there is a proposition *r* s.t. *p* contextually entails information about the likelihood that *r* is a partial answer to Q.

Then for questions (again, thanks to Simons):

(48) A question Q_1 is **RELEVANT** to question Q_2 just in case there is a proposition *p* such that any complete answer to Q_1 contextually entails information about the likelihood that *p* is a partial answer to Q_2.

Consider cases where *p* directly entails a partial answer to the QUD. 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. In this simple case, the new notion of RELEVANCE would agree 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 (as in the Holmes and Mastermind examples), the response itself is directly RELEVANT in the original sense, and adding it to the CG partially answers the QUD.

But under this revision, EMA-*p* would be RELEVANT even if only its prejacent provides a direct answer to the QUD. In such a case, both the entire EMA-*p* and the prejacent would be at-issue. Hence, in (26a-c) and (41a) the entire attitude proposition/EMA would be RELEVANT in the contexts given: though they don't directly address ?*p*, both bear on its likelihood, entailing that some agent has evidence consistent with its truth. When *?p* is the QUD, this is the only way we can understand such utterances to be RELEVANT. And since they are RELEVANT both (26a-c) and (41a) *and* their complement/prejacent *p* are at-issue. Hence Apt Response correctly predicts that we can respond anaphorically to either the entire attitude/modal proposition asserted or the non-asserted *p* (though in such a case, because both the attitude/EMA statement and its complement/prejacent are RELEVANT, typically a coda will be required to disambiguate the intended target for evaluation).

If we adopted some revision along these lines, it would still be useful to distinguish what is directly RELEVANT—those utterances that entail at least a partial answer to the QUD—from what is only indirectly so—those which only do so contextually, or only clarify what the live options are, etc.

I think it is generally wise to be cautious about revising RELEVANCE; we want the most constrained notion we can manage, to optimize its predictive force. But in the present case, I

[33] 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. For example, it becomes clear that we need to consider how to measure relative RELEVANCE: if *?p* is at-issue, there is a sense in which either of the answers *p* or *not-p* would be more informative than mere $\Diamond p$ or $\Diamond$*not-p*. Does this make the non-modal answers more RELEVANT than the modal ones? or is it merely a matter of Quantity? The rough revision in this section is just a place-holder for a deeper investigation, included here because it is RELEVANT for present purposes.

think it may be reasonable to explore the proposed modification, so that evidence for answers itself counts as RELEVANT.

7. Conclusions and prospects

Above I have argued that in responding to an utterance, agreeing is different than assessing. Agreement is about congruence of belief states (or cognitive states more generally), not the truth of a proposition expressed. And even in the case of assessing what someone has said, what's assessed isn't so much the proposition expressed as how it bears on the QUD. This conclusion, argued for above, has implications for a number of other proposals in the literature.

Recall that Swanson (2006) suggests that an epistemic modal claim can serve more than one kind of function:

(49) Functions of an EMA statement with prejacent *p*: [Swanson 2006]
- (i) it may assert something (proposing to add the proposition to the Common Ground/use it to reduce the CS);
- (ii) it may performatively raise the possibility that *p*; and
- (iii) it may advise an addressee on how to update her subjective probabilities as these bear on the likelihood that *p*.

I assume that in all the declarative uses considered above, the EMA-*p* statement is a Stalnakerian assertion, performing function (i). But when what's at-issue is *p* alone, (ii) the statement pragmatically (via RELEVANCE) raises the possibility that *p* (tacitly posing *?p*), and (iii) suggests a revision of the CG in such a way as to agree in the RELEVANT way with the belief state of the EMA-anchoring agent, which (because the CG is *common*) entails that cooperative addressees will update their own subjective probabilities with respect to *p* accordingly. Thus, we don't need speech act force 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. Moreover, by tying (ii) and what's at-issue (iii) to the QUD and RELEVANCE, we predict which of these functions will be evident in which contexts. And, finally, we predict the attested patterns of felicitous response and agreement, along with the attested interpretations of these speech acts, as a function of the particular context of utterance.

No account which attributes some of these effects to a judge parameter of interpretation (itself unconstrained by other contextual features) or to a speech act operator in the semantics can make the full range of correct predictions. Moreover, either of those other approaches requires ad hoc assumptions, indices, operators or ambiguities, whereas the present proposal depends almost entirely on independently motivated contextual factors and principles. The only novel suggestions here are that belief states are richer than just a set of propositions or the corresponding set of worlds in which all are true, and a proposal about how this richer notion is used to characterize what it is for two belief states to be congruent with respect to some proposition. But once we consider what it is for two agents to agree about the likelihood or relative plausibility of some proposition, I think it clear that some such notion will be required.

Prima facie, examples like (1), (4a,c), Moltmann's (29), and Rolf's response in (41c) have been taken to help make the case for Relativism: Alex and Rolf both make claims about the possibility that the explanation for Louise's absence is that she's busy with her dissertation, one promoting this as a possible explanation, the other denying; Rob's assertion likewise entails that that answer is not a possibility. The fact that each of them may truthfully make these claims based on their own solipsistic evidence doesn't seem to figure in their *disagreement*. Hence, Relativists assume that the disagreement is over a single non-sollipsistic modal proposition. But here we have explained these responses and their disagreement, based on independently motivated pragmatic principles, while assuming that the EMA propositions expressed by Alex and Rolf *are* solipsistically anchored.

The proposed account is superior to the Relativist approach in at least three other ways:

- Relativism cannot handle examples with bound variable anchors like that in (2): neither static contextualism nor relativism can account for these cases. See Roberts (2015a) for an indexical story about the anchoring of EMAs which handles these straightforwardly.
- Response patterns to Simons' parenthetical attitudes arguably display the same phenomenon as the EMAs, but no relativist account could plausibly address those: Doing that would have to involve some kind of monstrous operator shifting contexts of assessment (and their judges) under the scope of attitude predicates. Note that in these cases even when the response is to the complement alone the original speaker still isn't committed to its truth.
- Not only is the present account is independently motivated, it also has a wider scope of application than that of the Relativist approach to EMAs. See Roberts (2011), who argues, on empirical grounds, that projection of the prejacent of *only* is a function of the QUD; and see Simons et al. (2017), Beaver et al. (2017) on the role of at-issueness in the (non-)projection of factive complements. Hence, this general approach is motivated by more than just EMAs and parentheticals. Generally, more than one proposition may be made salient by an utterance, and which one we respond to is a function of what's at-issue, what's relevant to the purposes of the conversation.

Another account of interest is that of Ciardelli et al. (2009), couched in Inquisitive Semantics. They argue that EMAs are neither informative nor inquisitive—intuitively, they neither make an assertion nor pose a question, but instead have only attentive content: they draw attention to the proposition denoted by the prejacent. "[A]ttentive sentences. . .draw attention to possibilities that do not contribute to representing their informative or inquisitive content." (This general view was prefigured in Groenendijk, Stokhof & Veltman (1996), but worked out rather differently by Ciardelli et al.) I think the anaphoric response evidence argues against their contention that EMAs have no informative content, something like 'no proposition expressed': If not, then what is being assessed?

Finally, the idea that the QUD plays an important role in guiding the interpretation of an EMA is not new, though to my knowledge no one has previously used it in quite the way proposed here. Moss (2015) argues that the QUD plays a role in establishing the Modal Base with respect to which the domain of an epistemic modal is determined (Kratzer 1981), with truth conditional consequences when the same indicative modal clause in uttered in contextual minimal pairs. That proposal is just an extension of the now well-established general observation that in an alternative semantics the QUD (or topic) plays a role in domain restriction across a wide range of

operators (Rooth 1985,1992; Roberts 1989,1995,1996,2011; von Fintel 1994; Beaver & Clark 2008, *inter alia*). That the QUD should play a further role in the interpretation of epistemic modality, constraining what can be taken to be at-issue, has been independently argued for evidentials (Murray 2014). The proposal here just applies this observation to the case of EMAs, and extends its import, via Apt Response, to propose that we should understand patterns of attested response and their interpretation, as well as agreement in discourse, as essentially pragmatic.

Finally, the account I have offered here illustrates a theoretical issue of more general interest. It shows how a clear, independently motivated account of the pragmatics of discourse can help to explain the meanings of utterances in context, on the basis of relatively simple, non-ad hoc semantics, and with empirically superior results—predicting not only the range of meanings attested for individual sentences, but which meanings are likely to arise in which contexts of utterance. This is a central kind of factor to consider when we study the meaning of an expression; in fact, I would say it's the name of the game in the contemporary study of meaning.

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.

Anand, Pranav & Andrew Nevins (2004) Shifty operators in changing contexts. In K. Watanabe & R.B. Young (eds.) *Proceedings of SALT 14.* Ithaca, NY: CLC Publications.

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.

Baker, Carl (2012). Indexical contextualism and the challenges from disagreement. *Philosophical Studies* 157(1):107–123.

Barker, Chris (2002). The dynamics of vagueness. *Linguistics and Philosophy* 25(1), 1–36.

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, David & Brady Clark (2008) *Sense and Sensitivity: How Focus Determines Meaning*. Oxford: Wiley-Blackwell.

Beaver, David, Craige Roberts, Mandy Simons & Judith Tonhauser (2017) Questions Under Discussion: Where information structure meets projective content. *Annual Review of Linguistics* 3:265-84.

Büring, Daniel (2003) On D-Trees, Beans, and B-Accents. *Linguistics and Philosophy* 26(5):511-545.

Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen (2009) Attention! *Might* in inquisitive semantics. In E. Cormany, S. Ito & D. Lutz (eds.) *Proceedings of SALT 2009*, pp.91-108.
Cummins, Chris, Patricia Amaral & N. Katsos (2012) Experimental investigations of the typology of presupposition triggers. *Humana Mente* 23:1-16.
Davis, Christopher, Christopher Potts & Margaret Speas (2007) The pragmatic values of evidential sentences. In *Proceedings of SALT 17*, The Linguistic Society of America. DOI: http://dx.doi.org/10.3765/salt.v17i0.2966
Deal, Amy Rose (2013) Nez Perze embedded indexicals. In H. Greene (ed.) *Proceedings of SULA 7*. Amherst: GLSA.
DeRose, Keith (1991) Epistemic possibilities. *The Philosophical Review* 100(4):581-605.
Dowell, Janice (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
Dreier, J.l (2009) Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives* 23:79-110.
Egan, Andy (2007) Epistemic Modals, Relativism, and Assertion. *Philosophical Studies* 133:1–22.
Egan, Andy (2014) There's something funny about comedy: A case study in faultless disagreement. *Erkenntniss* 79:73-100.
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.
Faller, Martina (2006) Evidentiality above and below speech acts. http://semanticsarchive.net/Archive/GZiZjBhO/.
Farkas, Donka & Kim B. Bruce (2010) On reacting to assertions and polar questions. *Journal of Semantics* 27(1):81-118.
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.
Gibbard, Allan (1981) Two recent theories of conditionals. In W.L. Harper, R. Stalnaker & G. Pearce (eds.) *Ifs*. Dordrecht: D. Reidel, pp.21-247.
Ginzburg, Jonathan (2012) *The Interactive Stance*. Oxford University Press.
Grice, H.P. (1957) Meaning. *The Philosophical Review* 66 (1957).
Groenendijk, Jeroen & Martin Stokhof (1984) *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph.D. dissertation, University of Amsterdam.
Groenendijk, Jeroen, Martin Stokhof & Frank Veltman (1996) Coreference and modality. In Shalom Lappin (ed.) *The Handbook of Contemporary Semantic Theory*. Blackwell, 179-214.

Grosz, B. J. (1977) *The Representation and Use of Focus in Dialogue Understanding*. Technical Report No. 151, Artificial Intelligence Center, SRI International, Menlo Park, CA.
Hacking, Ian (1967) Possibility. *Philosophical Review* 76(2):143–68.
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.
Hooper, Joan B. (1975) On Assertive Predicates. In: Kimball, John P. (Ed.), *Syntax and Semantics* Volume 4, 91-124.
Huvenes, T. T. (2011). Varieties of disagreement and predicates of personal taste. *Australasian Journal of Philosophy*, 90(1), 167–181.
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.
Kneer, Marcus (nd.a) Epistemic modal claims: Data and 'data'
Kneer, Marcus (nd.b) Predicates of personal taste: empirical data.
Kölbel, Max (2014) Agreement and communication. *Erkenntniss* 79:101-120.
Kramer, Ruth & Kyle Rawlins (2011) Polarity particles: an ellipsis account. In *Proceedings of NELS 39*.
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.
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.
Lee, Jungmee (2011) *Evidentiality and its interaction with tense: evidence from Korean*. The Ohio State University PhD thesis.
López de Sa, Dan (2015) "Expressing disagreement: A presuppositional indexical contextualist relativist account". *Erkenntnis* 80:153-165.
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 1*. 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.
MacFarlane, John (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
Marques, Teresa (2014) Doxastic disagreement. *Erkenntnis* 80:121-142.
Martin, Scott (2014) Supplemental update. Ms., Nuance Communications. Submitted for review.
Matthewson, Lisa, Hotze Rullmann & Henry Davis (2007) Evidentials as epistemic modals: evidence from St'át'imcets. In Jeroen Van Craenenbroeck (ed.), *Linguistic variation yearbook 2007*, vol. 7. John Benjamins.
Moltmann, Friederike (2012) Two kinds of first-person-oriented content. *Synthese* 184:157-177.
Moss, Sarah (2015) On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics* 8.5:1-81.

Murray, Sarah E. (2010) *Evidentiality and the structure of speech acts*. Rutgers University PhD thesis. http://www.semanticsarchive.net/Archive/WViOGQxY/.
Murray, Sarah (2014) Varieties of update. *Semantics & Pragmatics* 7.2:1-53.
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.
Plunkett, David & Timothy Sundell (2013) Sworkin's interpretivism and the pragmatics of legal disputes. *Legal Theory* 19:242-281.
Roberts, Craige (1989) Modal Subordination and Pronominal Anaphora in Discourse. *Linguistics and Philosophy* 12.6:683-721. Reprinted in Javier Gutierrez-Rexach (ed.) *Semantics: Critical concepts in linguistics*, Routledge, 2003.
Roberts, Craige (1995) Domain Selection in Dynamic Semantics. In Emmon Bach, Eloise Jelinek, Angelika Kratzer, and Barbara H. Partee (eds.) *Quantification in Natural Languages*, Kluwer, Dordrecht.
Roberts, Craige (1996/2012) Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. In Jae Hak Yoon and Andreas Kathol (eds.) *Ohio State University Working Papers in Linguistics* Volume 49. Reprinted in the 1998 version with a new Afterword in *Semantics and Pragmatics*, Volume 5, 2012.
Roberts, Craige (2004) Context in dynamic interpretation. In Laurence Horn and Gregory Ward (eds.) *Handbook of Contemporary Pragmatic Theory*, Blackwell, 197-220.
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, Lattvia.
Roberts, Craige (2015a) The Character of epistemic modality: Evidentiality, indexicality, and what's at-issue. Ms., The Ohio State University and Central European University.
Roberts, Craige (2015b) Plans and imperatives: A semantics and pragmatics for imperative mood. *Proceedings of the Amsterdam Colloquium 2015*. ILLC, University of Amsterdam.
Roberts, Craige (2016) Coherence and anaphora. Ms. of a talk delivered at the NASSLLI 2016 workshop on Anaphora and Coherence, Rutgers, July, 2016.
Rooth, Mats (1985) *Association with Focus*. Ph.D. dissertation, University of Massachusetts, Amherst.
Rooth, Mats (1992) A theory of focus interpretation. *Natural Language Semantics*, 1(1):75-116.
Sadock, Jerrold M. & Arnold M. Zwicky (1985) Speech act distinctions in syntax. In T. Shopen (ed.) *Language Typology and Syntactic Description*. Cambridge University Press, pp.155-196.
Sandt, Rob van der & Emar Maier (2003) Denials in discourse. Talk at the Michigan Philosophy and Linguistics Workshop. Ms., University of Nijmegen, Netherlands. Ms. available at https://sites.google.com/site/emarmaier/publications.
Schlenker, Philippe (2003) A plea for monsters. *Linguistics and Philosophy* 26:29-120.
Schwarzschild, Roger (1999) GIVENness, AvoidF and other Constraints on the placement of accent. *Natural Language Semantics* 7.2:141-177.
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 (2011) 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, Craige Roberts & Judith Tonhauser (2017) The Best Question: Explaining the projection behavior of factives. *Discourse Processes* 54.3:187-206, a special issue on the Question In Discussion. Available open-access at: http://www.tandfonline.com/doi/pdf/10.1080/0163853X.2016.1150660?needAccess=true

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. (1978) Assertion. In P. Cole (ed.), *Pragmatics*, New York: New York Academic Press, Vol. 9, pp. 315–32. Reprinted in Stalnaker 1999: 78–95.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.

Stojanovic, Isadora (2007) Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy* 30:691-706.

Sundell, Timothy (2011) Disagreements about taste. *Philosophical Studies* 155:267–288.

Swanson, Eric (2006) *Interactions with Context*. Ph.D. dissertation, MIT.

Syrett, Kristen & Todor Koev (2015) Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics* 32, 525-577.

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.

Tonhauser, Judith, David Beaver & Judith Degen (2017) The information-structural source of projection variability. Ms, OSU/UTexas-Austin/Stanford.

Tonhauser, Judith, David Beaver, Judith Degen, Marie-Catherine de Marneffe, Craige Roberts & Mandy Simons (2015) Negated evaluative adjective sentences: What projects, and why? Talk at Experimental Pragmatics 2015, University of Chicago, July 16-18, 2015.

Weatherson, Brian 2009. Conditionals and Indexical Relativism, Synthese 166/2: 333‑57.

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.