

The semantics of attitude markers and other illocutionary content

August 2016

Word count: 14,921

Abstract

Intuitively, attitude markers like *alas* encode not-at-issue information about the speaker's attitude towards the content of the sentence it occurs in (Frege, 1956; Vanderveken, 1990). Using attitude markers as a case study, I argue that there is in fact a difference between this type of content and canonical not-at-issue content, like that encoded in utterance modifiers like *frankly* or evidential adverbs like *apparently*. In contrast to the latter, attitude markers are only compatible with declarative mood; are uniquely restricted in terms of the content they apply to; and can result in Moore's Paradox. I conclude that the contribution of attitude markers should thus be treated as 'illocutionary content', on par with the sincerity conditions encoded in illocutionary mood. I present a formal analysis of illocutionary content in which it differs from other not-at-issue content in restricting the speaker's Discourse Commitments (Gunlogson, 2001) rather than the Common Ground.

1 Introduction

This paper lies at the intersection of two traditions of meaning distinction. The first – characterized by Speech Act Theorists like Stenius (1967); Searle (1969) but also recent dynamic adaptations like Farkas and Bruce (2010); Murray (2014) – distinguishes between the denotation of a sentence (e.g. its propositional content) and how the denotation affects or is applied to the context (e.g. the contribution of its illocutionary mood). The second – characterized by Potts (2005) and Simons et al. (2010) but also recent dynamic adaptations like Murray (2010); AnderBois et al. (2010) – distinguishes between the at-issue or non-projective content of a sentence and its not-at-issue or projective content (e.g. conventional implicatures and some presuppositions).

The goal of this paper is to argue that both of these distinctions are useful. There are two types of content that qualify as not-at-issue: the first type, what we canonically think of as not-at-issue content (encoded in appositives, Pottsian

expressives like *bastard*, evidential adverbs, and utterance modifiers like *frankly*), is part of the descriptive content of the sentence. The second type, which I refer to as ‘illocutionary content,’ is also semantically encoded and not-at-issue, but it instead pertains to how the denotation affects or is applied to the context.¹

Attitude markers, the empirical focus of this paper, are morphemes and prosody that mark a speaker’s attitude towards some descriptive content. Moreover, the information they contribute is not-at-issue and, as I will argue, is illocutionary content specifically. The sentence in (1-a) does not contain an attitude marker; the sentences in (1-b) and (1-c) do.

- (1) a. John lost the race.
- b. Alas, John lost the race.
- c. (Wow,) John lost the race!

An utterance of (1-b), with the attitude marker *alas*, indicates that the speaker is disappointed that John lost the race. (1-c), uttered with exclamation intonation (signified by the exclamation point and encouraged by particles like *wow*), indicates that the speaker is surprised that John lost the race.

The semantic contribution of these attitude markers is decidedly not-at-issue: their meaning cannot be targeted by truth-conditional operators, denied in discourse, or used to address the Question Under Discussion, as I will demonstrate below. But attitude markers differ in some notable ways from canonical encoders of not-at-issue content, like expressives and appositives: they must take narrow scope in certain contexts, for instance, and they can only occur with certain illocutionary moods.

In this paper, I’ll argue that attitude markers form a natural subclass of encoders of not-at-issue content in terms of the kind of meaning they encode (i.e. the speaker’s attitude) and their linguistic behavior. I will argue that the former partially determines the latter: attitude markers, in contrast to e.g. expressives and appositives, behave the way they do because a) they (necessarily) target propositional content (in contrast to sub-propositional content); and b) the information they encode pertains to the speaker’s attitude. Consequently, I’ll argue, attitude markers quite naturally encode their meaning at an illocutionary, rather than descriptive, level: they contribute to the speaker’s Discourse Commitments (Gunlogson, 2001), rather than the Common Ground. While others (Searle and Vanderveken, 1985; Vanderveken, 1990) have proposed a similar treatment, I attempt to better substantiate the claim, and formulate the analysis in a dynamic update semantics, which allows for a particularly natural compositional account of illocutionary mood and related content.

In the next section, I will better delineate the class of attitude markers by providing several examples and descriptive generalizations. In §3, I will show how they differ from apparently similar not-at-issue phenomena like evidentials,

¹This distinction loosely parallels one made in Kaplan (1997) (and underscored in Kratzer, 1999): “A descriptive is an expression which describes something which either is or is not the case. . . . [A]n expressive. . . expresses or displays something which either is or is not the case”. I use the term “illocutionary” instead of “expressive”, as the latter is confusing due to the fact that not all of Potts’ (2005) expressives are in fact expressive in this Kaplanian sense.

expressives, and discourse particles.

Two terminological notes: following many, exemplified by Hausser (1980), I will use the term **illocutionary mood** to refer to the morphosyntactic category marking sentence type. And I will use the term **illocutionary force** to label that which results in the speech acts these sentences are used in. To illustrate: Some languages morphologically mark imperative mood; because grammar underdetermines speech acts, imperative mood is consistent with a variety of illocutionary forces, including commands, suggestions, advice, etc.²

I will use the term **content** very broadly, to refer to any meaning that is lexically or prosodically encoded. As suggested above, I will be making two independent distinctions: that between **at-issue** and **not-at-issue content**, as decided by the standard tests of projection outside of a truth-conditional operator, deniability in discourse, and ability to address the Question Under Discussion. And also between **descriptive** and **illocutionary content**, loosely equivalent to Kaplan’s descriptive and expressive distinction. Illocutionary content is semantically encoded content pertaining to how the sentence affects or is applied to the context. This typology is illustrated in Figure 1.

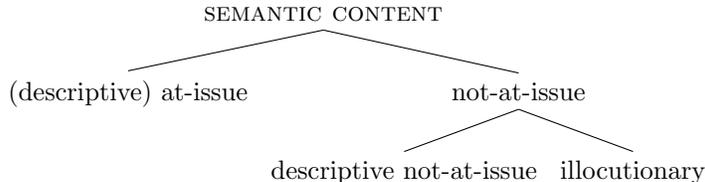


Figure 1: Content distinctions

I will introduce tests for this second distinction; it will turn out that all at-issue content is descriptive, and all illocutionary content is not-at-issue.

While the narrow goal of this paper is a characterization and explanatory compositional account of attitude markers, I believe the discussion here sheds light on a few larger issues. It addresses and attempts to delineate – for the first time, in my knowledge – the oft-implied intuition that there is a level of illocutionary content separate from canonical not-at-issue content like conventional implicature (see Frege, 1956; Potts, 2003b; Rett and Murray, 2013, for mentions of such a distinction). And it brings an interesting perspective to a recent debate involving cross-linguistic semantic variation in evidentials (Murray, 2010; Matthewson, 2011; Faller, 2014).

2 A profile of attitude markers

Attitude markers encode the speaker’s attitude towards the descriptive meaning of an utterance in backgrounded, not-at-issue content. To illustrate this, I focus initially on the morphemes *alas* and *fortunately* in English because they are

²Thanks to anonymous *Semantics & Pragmatics* reviewers for pushing me on this point.

lexical (as opposed to prosodic) and unambiguously target propositions. Later in this section, I will also discuss exclamation intonation and expressives like *damn* which may or may not target propositions.

I take the minimal pair in (2), repeated from (1), to illustrate the semantic contribution of an attitude marker (in this case, *alas*):

- (2) a. John lost the race.
b. Alas, John lost the race.
- (3) Fortunately, John lost the race.

While both utterances amount to an assertion that John lost the race, in (2-b) the speaker additionally conveys that she is upset or dismayed that John lost the race. *Fortunately* generally behaves like the antonym of *alas* (3); it's used to express that the speaker is pleased or relieved at the descriptive content of the utterance.

Evidence that attitude markers do not contribute to the at-issue content of the utterance; their content cannot be targeted by truth-conditional operators (4), and it cannot be denied in discourse (5).³

- (4) a. Alas, John did not lose the race.
b. Alas, it is not the case that John lost the race.
c. It is not the case that John lost the race, alas.
- (5) A: Alas, John lost the race.
B: That's not true, he won!
B':#That's not true, you're glad he did!

(4) shows different ways of negating the sentence in (2-b); none can negate the contribution of *alas*. These utterances cannot be used to express that the speaker does not regret that John did not lose the race. In (5), Speaker B's protest that Speaker A's utterance is not true can be justified by the claim that John won; in contrast, the B' protest cannot be justified by the claim that Speaker A was glad that John lost the race.

In contrast, embedding verbs like *be disappointed* or *be surprised* are not attitude markers, because their contribution tends to be at-issue, and is in every case targetable by truth-conditional operators (e.g. *I am not surprised that John won the race*). In the rest of this section, I will provide some additional examples of attitude markers, all of which pattern like *alas* and *fortunately* in tests for not-at-issueness. In the following section, I will show how attitude markers form a distinct class within the realm of not-at-issue markers.

Instead of disappointment or relief, many attitude markers indicate that the speaker is surprised by (or had not expected) the propositional content of the utterance. This phenomenon is, in some traditions, referred to as 'mirativity' DeLancey (1997, 2001).

³The English judgments presented in this paper are those of the author and all of the four other native speakers of English informally consulted for judgments. I will note disagreement where it has arisen.

In English, speaker surprise can be marked intonationally, by a prosodic attitude marker. This is illustrated by the difference between the assertion in (6-a) and the exclamation in (6-b) (Sadock, 1974; Cruttenden, 1986; Michaelis, 2001; Merin and Nikolaeva, 2008). I take the exclamation point in these examples to model a particular prosody or intonation in English: a steady Rise, abrupt Fall contour (Cruttenden, 1986) plus features of emphasis like lengthening (Bartels, 1999). This intonation is brought out especially well by discourse particles like *wow*, although such discourse particles should not be confused with the attitude marker (the intonation) itself: they are optional in exclamatives, and can occur on their own, without descriptive content.

- (6) a. John arrived on time.
 b. (Wow,) John arrived on time!

The difference in meaning between (6-a) and (6-b) is one of speaker attitude: both utterances convey the same descriptive content, but an utterance of (6-b) additionally expresses that the speaker is surprised by (or had not expected) the descriptive content.

Exclamation intonation behaves like *alas* in that its content seems to be not-at-issue. It can't be denied in discourse (7), and it cannot be targeted by negation (8).

- (7) A: (Wow,) John lost the race!
 B: That's not true, he won.
 B':#That's not true, you knew he would lose.
- (8) (Wow,) John did not lose the race!

Like (4), (8) means 'The speaker is surprised that John did not lose the race.' It cannot mean 'The speaker is not surprised that John lost the race.'

While English encodes speaker surprise or mirativity in intonation, other languages encode the meaning lexically. In Finnish, for example, the sentence particle *-pä* expresses speaker surprise (Karlsson, 1999, 20). Like the pair in (6), the pair in (9) differ only in that (9-b) additionally encodes that the speaker finds the propositional content (that there are lots of flowers) surprising. They do not differ in intonation.⁴

- (9) a. Täällä on paljon kukk-ia.
 here be-3RD.SG a.lot flower-PRT.INDF.PL
 'There are lots of flowers here.'
- b. Täällä-pä on paljon kukk-ia.
 here-PA be-3RD.SG a.lot flower-PRT.INDF.PL
 '(Wow,) There are lots of flowers here!'

Like exclamation intonation, the content encoded by *pä* cannot be directly denied in discourse, and cannot be targeted by negation.

⁴Thanks to Peter Sutton (p.c.) for drawing my attention to *pä*, and to Tuomo Tiisala (p.c.) for his judgments.

Wu (2008) reports two sentential adverbs in Mandarin that seem to count as speaker attitude markers. As shown in (10), *jingran* is a mirative marker, and the adverb *guoran* seems to be its antonym: it is used to express that the asserted content was expected by the speaker.

- (10) Zhangsan guoran /jingran lai le.
 Zhangsan GUORAN /JINGRAN come PST
 ‘Zhangsan came (as expected/not expected by the speaker).’

Finally, as detailed in Rett and Murray (2013) and elsewhere, there is a robust crosslinguistic tendency for indirect evidential markers to double as mirative markers. I’ll briefly introduce the phenomenon of evidentials and then illustrate mirative evidentials from Tsafiki, a Barbacoan language spoken in Ecuador, as reported in Dickinson (2000).

Tsafiki is an evidential language, which means that all grammatical sentences contain an evidential marker that specifies the type of evidence for their descriptive content (Aikhenvald, 2004). Tsafiki’s is a three-way evidential system; it distinguishes between direct physical evidence (11-a), information inferred from direct physical evidence (11-b), and information inferred from general knowledge (11-c) (from Dickinson, 2000, 407–8).⁵

- (11) a. Manuel ano fi-e.
 M food eat-DECL
 ‘Manuel ate.’ (The speaker saw him.)
 b. Manuel anno fi-nu-e.
 M food eat-IND-DECL
 ‘Manuel ate.’ (The speaker sees the dirty dishes.)
 c. Manuel anno fi-n-ki-e.
 M food eat-NOM-INF-DECL
 ‘Manuel must have eaten.’ (He always eats at 8:00; it’s now 9:00.)

However, in certain contexts, the indirect evidential *nu* marks mirativity instead of indirect evidence. Dickinson (p411) describes (12) as ambiguous.

- (12) Moto jo-nu-e.
 motorcycle be-IND-DECL
 ‘It is a motorcycle.’ (The speaker hears a motor.)
 ‘It’s a motorcycle!’

In a context in which the speaker has indirect evidence for the proposition – for instance, that she hears rather than sees a motorcycle – the indirect evidential in (12) is licensed. However, it is also licensed in contexts in which the speaker sees the motorcycle, i.e. has direct evidence for the motorcycle, in which case that same evidential functions instead as a mirativity marker, in which case (12) conveys that the speaker is surprised that it is a motorcycle.

⁵DECL labels the declarative marker, which encodes assertive illocutionary mood. I’ve labelled the second evidential IND for ‘indirect’ and the third INF for ‘inferential’.

The polysemy illustrated in (12) – the repurposing of indirect evidentials as mirative markers – happens across languages and language families: in Turkish (Slobin and Aksu, 1982), Tibetan (DeLancey, 1997), and Cheyenne (Rett and Murray, 2013). Important here is the observation that, when these polysemous morphemes receive a mirative interpretation, they are acting as attitude markers.

In the next section, I will argue that attitude markers behave differently in principled ways from other encoders of not-at-issue content, including evidentials.⁶

3 What distinguishes attitude markers

I’ve defined attitude markers as linguistic elements that mark a speaker’s attitude towards the descriptive content of the utterance. This definition excludes a variety of canonical encoders of not-at-issue content. I’ll first explain the contrasts I have in mind, and then show that attitude markers do in fact behave as a distinct subclass of encoders of not-at-issue content.

3.1 What attitude markers aren’t

There are a number of other encoders of not-at-issue content: most recognizably presupposition triggers; Pottsian expressives like *bastard*; and appositives. These are illustrated in (13); (14) demonstrate that their meanings, too, cannot be targeted by truth-conditional operators like negation.

- (13) a. John’s sisters live in Melbourne. *presupposition trigger*
 b. John met with that bastard Bill. *Pottsian expressive*
 c. John, an avid cyclist, won the race. *appositive*
- (14) a. It’s not the case that John’s sisters live in Melbourne.
not negated: John has sisters
 b. It’s not the case that John met with that bastard Bill.
not negated: The speaker thinks that Bill is a bastard
 c. It’s not the case that John, an avid cyclist, won the race.
not negated: John is an avid cyclist

Non-presupposed, lexically encoded not-at-issue content is generally characterized as conventional implicature (Potts 2005, though see Bach 1999). Conventional implicature is typically analyzed semantically in one of two ways: statically, on a different tier from at-issue content (Potts, 2003a,b, 2005, 2007; Gutzmann, 2015); or dynamically, as an automatic common-ground update (Murray, 2010, 2011, 2014; AnderBois et al., 2010, 2013). I’ll ultimately argue that these approaches do not appropriately account for the idiosyncratic behavior of attitude markers.

⁶Correspondingly, I will strengthen the above claim to a biconditional: *only when* the polysemous mirative evidentials receive a mirative interpretation do they qualify as attitude markers.

There are a number of encoders of not-at-issue content that don't count as attitude markers given this definition. Appositives, and many expressives (in e.g. *that bastard John*) sometimes reflect the speaker's attitude, but it is not an attitude directed at the descriptive content of the utterance (i.e. a proposition), but rather something sub-propositional, like an individual. An exception is *damn*, which can modify individuals (e.g. *that damn postman*) or can be used to mark the speaker's dismay at a proposition (e.g. *Damn, John lost the race!*). When it is used in the latter sense, *damn* qualifies as an attitude marker.⁷

And while many other types of not-at-issue content target propositions, they do not encode the speaker's attitude towards it. Evidentials encode the speaker's type of evidence for the descriptive content of an utterance, and speaker certainty markers like *of course* encode the speaker's level of credence in it (see also Ettinger and Malamud, 2015). Utterance modifiers like *frankly* appear to modify the speech act itself, causing Potts (2003b) to analyze them (along with Japanese performative honorifics) as denoting not-at-issue content associated with a null verb *utter*.

The phenomenon of discourse particles (e.g. English *too, even, indeed*, German *toch, doch*) is also relevant; however, as the label 'discourse particle' describes a morphosyntactic category, it refers to a semantically heterogeneous class (see Waltereit, 2001, for a related discussion). Zimmermann (2011) defines the class of discourse particles as follows (p2012): "Discourse particles in the narrow sense are used in order to organize the discourse by expressing the speaker's epistemic attitude towards the propositional content of an utterance, or to express a speaker's assumptions about the epistemic states of his or her interlocutors concerning a particular proposition." Particles that satisfy this first disjunct, like the Finnish *pää*, qualify as attitude markers according to the present definition; those that satisfy the second do not. And in fact, Zimmermann observes that discourse particles can differ in their behavior in just the way I will claim that attitude markers differ from other encoders of not-at-issue content: "While some discourse particles add extra meaning to the descriptive content in form of a presupposition or conventional implicature..., others operate more directly on the descriptive and illocutionary meaning, e.g. by changing the strength of speaker or hearer commitment towards the proposition expressed" (p2015-6). This is the topic of the next subsection.

3.2 How attitude markers differ

In the previous subsection, I delineated the class of attitude markers based on their characterization as not-at-issue encoders of the speaker's attitude towards some content of an utterance. In this section, I'll show that this characterization appropriately identifies a natural subclass of not-at-issue content: illocutionary content, which is not compatible with all illocutionary moods; has a relatively restricted interpretation in conditionals; and can result in Moore's Paradox. I'll discuss these in turn.

⁷I am indebted to Kai von Fintel (p.c.) for making this point.

3.2.1 Acceptability across illocutionary mood types

The previous section discussed the ability to be targeted by truth-conditional operators like negation as a test for not-at-issue content. Another test is whether or not the content persists in questions. The questions in (15) carry the presuppositions and implications that their declarative counterparts in (13) do, but the entailments of those sentences do not survive.

- (15) a. Where do John’s sisters live?
 b. Where did John meet with that bastard Bill?
 c. When did John, an avid cyclist, win the race?

Attitude markers, in contrast, are unacceptable in constituent questions, and, for the most part (as I’ll demonstrate shortly), in polar questions. This general unacceptability in questions seems to be a unifying characteristic of attitude markers, and it is the starkest difference between attitude markers and other encoders of not-at-issue content.

- (16) a. *Alas/Fortunately/Unfortunately, who lost the race?
 b. *Who lost the race, alas/fortunately/unfortunately?

English exclamation intonation is unacceptable in questions (although this is plausibly for phonological reasons, since questions in English are marked in part prosodically). But the Finnish mirativity marker *pä* is also unacceptable in questions, across the board, despite being a lexical attitude marker.

It is also true of mirative evidentials in Cheyenne. Recall that mirative evidentials mark evidentiality in some contexts and mirativity (or speaker surprise) in others. In the latter cases, they count as attitude markers. And, as shown in Rett and Murray (2013), this difference in interpretation affects their ability to occur in questions. When the Cheyenne mirative evidential (glossed as NAR for ‘narrative’) occurs in questions, as in (17), it can only receive an evidential interpretation, not a mirative interpretation. (The unavailability of the intended interpretation in (17-b) is marked by %.)

- (17) a. Mó=é-x-hó’ táhevá-**hoo’o** Aénohe?
y/n=3-REM.PST-win-NAR.3SG Hawk
 ‘Given the stories you heard, did Hawk win?’
 b. %Mó=é-hó’ táhevá-**hoo’o** Aénohe?
y/n=3-win-NAR.3SG Hawk
 Intended: ‘Given your surprise, did Hawk win?’ / ‘Did Hawk really win?!’

Attitude markers like *fortunately* and *unfortunately* are unacceptable in polar questions as well as constituent questions, as demonstrated in (18).

- (18) a. *(Un)fortunately, did John lose the race?
 b. *Did John lose the race, (un)fortunately?

However, *alas* is relatively acceptable in polar questions, as (19) demonstrates.

(19) Alas, did John lose the race?

It's worth noting that, in (19), the attitude marker doesn't target the question: it does not convey that the speaker is dismayed at whether John lost the race. In (19) *alas* instead targets the bias of the polar question, namely the proposition that John did lose the race (Büring and Gunlogson, 2000; Sudo, 2013). In §4 I'll argue that attitude markers are restricted across illocutionary moods because they range over utterances with one (and only one) salient proposition. While there is room in this account for attitude markers to be anaphoric to one salient proposition in a set, as appears to be the case in (19), it cannot explain the contrast between *alas* and *unfortunately* (exemplified in (18) and (19)).⁸

In addition to being generally unacceptable in questions, there seems to be a larger prohibition of attitude markers in utterances with non-declarative mood: they are unacceptable, for instance, in imperatives. This is true for Finnish *pä* (and English exclamation intonation, again for plausibly phonological reasons); I illustrate it below for some others.

(20) a. #Alas/Unfortunately, clean up your room!
b. #Fortunately, get a job!

In sum, while not-at-issue content encoded in phenomena like expressives and appositives projects outside of questions (15), attitude markers tend to be unacceptable in questions, and with non-declarative mood in general. The ability to predict this difference is thus an important desideratum for a semantic theory of attitude markers.

3.2.2 Differences in scope

As discussed above, there are a number of canonical encoders of not-at-issue content that can range over propositions: utterance modifiers like *frankly*; evidentials like *apparently*; and some Pottsian expressives like *damn*. In this subsection, I will show that attitude markers differ from these other elements in their scope, by which I mean the propositions they apply to. In some constructions, like the antecedents of conditionals, their scope is relatively restricted.

Canonical encoders of not-at-issue content can occur in the antecedents of conditionals as well as sentence-initially.⁹ When they do, they range over the

⁸It's been suggested to me that lexical differences between attitude markers are best attributed to diachronic factors; in particular, that *alas* is more distantly related to its adjectival or adverbial roots than *unfortunately* (e.g. *It was unfortunate/*alas that John lost the race*).

⁹Interestingly, attitude markers must syntactically precede these other sentential modifiers, demonstrated by the contrast in (21) (Cinque, 1999).

(21) a. Alas, apparently John died.
b. *Apparently, alas John died.

Neither seems to be able to apply to the other's content, however; (21-a) cannot mean the speaker is disappointed that it's apparent (as opposed to unapparent, or clear) that John died.

conditional as a whole. This is illustrated in (22) for the utterance modifier *frankly* and in (23) and (24) for the evidentials *apparently* and *allegedly*.¹⁰

- (22) a. Frankly, if the mayor is convicted, she must resign from office.
b. If, frankly, the mayor is convicted, she must resign from office.
- (23) a. Apparently, if the mayor is convicted, she must resign from office.
b. If, apparently, the mayor is convicted, she must resign from office.
- (24) a. Allegedly, if they didn't follow protocol, they'll be fired.
b. If, allegedly, they didn't follow protocol, they'll be fired.

In other words: when canonical encoders of not-at-issue content are in the antecedent of conditionals, as they are in the (b) sentences, they cannot range over the antecedent itself, but must range over the conditional as a whole. This claim is illustrated more definitively in (25-b), in which the antecedent but not the conditional as a whole is compatible with indirect evidence. The resulting conditionals are unacceptable, regardless of the position of the evidential.

- (25) a. #Apparently, if the mayor is convicted, I will run for office.
b. #If, apparently, the mayor is convicted, I will run for office.

To summarize, canonical encoders of not-at-issue content project globally in conditionals: regardless of their syntactic position in the conditional, they range over the conditional as a whole.

Lexically encoded attitude markers have the same distribution in conditionals, but their interpretation in antecedents is, in contrast, local. This is illustrated below with *alas*.

- (26) a. Alas, if the mayor is convicted, she must resign from office.
b. If, alas, the mayor is convicted, she must resign from office.

In (26-a), *alas* can range over the entire conditional, as do the not-at-issue encoders in (22)–(24). But in (26-b), when *alas* is in the antecedent of the conditional, it can only range over the antecedent, not the meaning of the conditional as a whole.

This point is illustrated more definitively by the contrast in (27), a conditional in which the consequent is a proposition the speaker would (in a neutral context) not be dismayed by.

- (27) a. #Alas, if the mayor is convicted, at least we'll have the chance of getting a better one.
b. If, alas, the mayor is convicted, at least we'll have a chance of getting a better one.

Because the content encoded in *alas* is incompatible with the chance of getting a better mayor, and because the sentence-initial *alas* in (27-a) necessarily ranges

¹⁰I do not address the behavior of these elements embedded in the consequent, as it is too hard to tell their scope in these positions.

over the conditional as a whole, the conditional in (27-a) is unacceptable. In contrast, since *alas* only ranges over the antecedent when it is embedded in the antecedent, (27-b) is acceptable. This pair forms a direct contrast with the pair in (25-a) and (25-b).

Other lexical attitude markers, like *fortunately* and *unfortunately*, behave the same way as *alas* in conditional antecedents. However, it's worthwhile noting that the expressive *damn*, which I have classified as an attitude marker when it ranges over propositions, cannot be embedded anywhere in conditionals.

- (28) a. Damn, if the mayor is convicted, she'll have to resign!
 b. *If, damn, the mayor is convicted, she'll have to resign!
 c. *If the mayor is convicted, she will, damn, have to resign!

I'll argue in §4 that this variation across attitude markers can be explained in terms of the requirements they place on the propositional argument to which they are anaphoric.

Conditionals aren't the only construction in which attitude markers must scope locally; the conjunctions in (29) and the appositives in (30) illustrate the same paradigm as the conditionals in (26).

- (29) a. Alas, John and Mary attended the party.
 b. John and, alas, Mary attended the party.
- (30) a. Alas John, who won the lottery, lost the race.
 b. John, who alas won the lottery, lost the race.

In particular, the most natural interpretation of (29-a) is that the speaker is disappointed that John and Mary attended the party, while the most natural interpretation of (29-b) is that the speaker is disappointed that Mary (but not necessarily John) attended.

Notice here that the semantic complement of *alas* is still propositional: the attitude marker isn't merely ranging over the denotation of *Mary*, as (29-b) is consistent with a context in which the speaker loves Mary (but is just, for other reasons, dismayed that she attended the party). The example in (30-b) additionally illustrates that, in certain configurations, attitude markers can apply to an utterance's not-at-issue content.

It's possible that the semantic complements of attitude markers in these sentence-medial contexts are constructed via ellipsis. But it's more plausible that they come about via the same mechanisms as focus alternatives do, because one further difference between attitude markers and canonical encoders of not-at-issue content is that attitude markers are focus-sensitive.

- (31) a. Alas, JOHN_F got a dog.
 b. Alas, John got a DOG_F.
- (32) a. (Wow,) JOHN_F got a dog!
 b. (Wow,) John got a DOG_F!

I use all caps and a subscript *F* to represent focus-marking. In the (a)

sentences above, the attitude markers apply to the proposition that John and not someone else got a dog. In the (b) sentences, they apply to the proposition that John got a dog and not something else. The meaning of Finnish *pä* can vary in a similar way as in (32), but it does so by varying the relative position of *pä* (Karttunen and Kay, 1985; Vallduví and Vilkuna, 1998).

- (33) a. Minä-pä soitan nyt isoäidille.
 I-PA call-1SG-NON.PST now grandmother
 ‘I_F am calling grandmother now.’
 b. Isoäidille-pä soitan nyt.
 grandmother-PA call-1SG-NON.PST now
 ‘I am calling GRANDMOTHER_F now.’

The examples in (33) additionally differ in the word order of the sentence, as Finnish has relatively free word order, but *pä* must occur in the highest functional projection of the clause (Zimmermann, 2011), and is therefore typically suffixed to the first phrase in the sentence.

3.2.3 Moore’s Paradox

There is a final way in which attitude markers differ from canonical encoders from not-at-issue content: denying the content of attitude markers results in Moore’s Paradox, while denying the content of other not-at-issue meaning results in something more like a contradiction. This test was borrowed from Murray (2010)’s work on Cheyenne – which is relevant to the distinction here – and adapted slightly.

Recall that mirative evidentials, like the Cheyenne mirative evidential, receive an evidential interpretation in some contexts and a mirative interpretation in others. When they are interpreted miratively, they count as attitude markers.

Moore’s Paradox, exemplified in (34), results from the assertion of a proposition coupled with the speaker’s denial that she believes the proposition.

- (34) #It’s raining, but I don’t believe it’s raining.

The source of this unacceptability has been the cause of a great deal of debate; in the next section, I will side with Searle (1969) in arguing that Moore’s Paradox occurs when the second conjunct denies a sincerity condition encoded in the illocutionary content of the first.

Murray reported that the evidential and mirative uses of Cheyenne mirative evidentials differed in Moorean constructions. She presented consultants with a pair of conjoined sentences; the first conjunct in each contained the mirative evidential, the second conjunct denied the content of the mirative evidential. In the first sentence, in (35), the mirative evidential had an evidential interpretation (in Cheyenne, the reportative). In the second sentence, (36), the mirative evidential had a mirative interpretation.

- (35) #₁É-hó’tàheva-séstse Aénohe naa oha hovánee’e
 3-win-RPT.3SG Hawk but nobody

é-sáa-nè-hé-he-∅.

3-NEG-that-say-MOD_A-DIR

Intended: ‘Hawk won, it’s said, but nobody said that.’

- (36) #É-hoo’kóhó-**neho!** Ná-nèšè-héne’ena tsé-to’sè-hešè-hoo’koho.
3-rain-nar.sg.inan 1-continue-know.s.t CNJ-going.to-how-rain
Intended: ‘It’s raining! ... # I *knew* it was going to rain.’

Both sentences sounded unacceptable to Murray’s consultants. But they reported a difference in the unacceptability. They claimed that the mirative Moore’s Paradox sentence in (36) was infelicitous (marked with #), but the evidential Moore’s Paradox sentence in (35) was **contradictory** (marked #_⊥). This former judgment was the same one the consultants reported for traditional Moorean sentences like (34).

Murray interpreted this as evidence that mirative content differed in kind from evidential content, and in particular that the former patterned with the sincerity conditions of an utterance (and were thus better characterized as illocutionary content than the evidential interpretation of Cheyenne mirative evidentials). I will adopt the same conclusion.

It’s possible, however, that native speaker intuition cannot detect differences in kinds of infelicity at this level of subtlety. When I consulted native English speakers about this distinction, presented in the form of a sorting task, about half were unable to recognize a difference between the type of felicity in any of the three sentences below: (37-a) the classic Moore’s Paradox; (37-b) a version with canonical not-at-issue content (an evidential adverb); and (37-c) a version with attitude markers.

- (37) a. #It’s raining, but I don’t believe it’s raining.
b. #Allegedly, John lost the race, but no one alleged he did.
c. #Alas, John lost the race, but I’m not disappointed he did.

It is worthwhile noting, however, that those consultants who did report differences in their judgments reported that the attitude marker sentence in (37-c) patterned with the classical Moorean sentence in (37-a) to the exclusion of the evidential sentence in (37-b), which is consistent with Murray’s findings.

An apparently clearer test involves the embedding of Moorean sentences in certain epistemic contexts (Yalcin, 2007). Standard Moorean sentences like (37-a) become acceptable when embedded in the antecedent of a conditional or under the verb *suppose*, as in (38) (although see Roberts 2015 for a dissenting view).

- (38) a. Suppose it is raining, but I do not believe that it is raining.
b. If it is raining but I don’t believe it, then there is something I do not believe.

Since the previous subsection demonstrated that attitude markers and other encoders of not-at-issue content differ in their scope in the antecedents of conditionals, I will restrict myself to the *suppose* test.

- (39) a. Suppose that, alas, John lost the race, but that I'm not disappointed he did.
 b. #Suppose that, allegedly, John lost the race, but no one alleged that he did.

The sentences in (39) illustrate a contrast, arguably similar to the one reported by Murray for Cheyenne: denying the content encoded in attitude markers like *alas* results in Moore's Paradox, which is obviated under *suppose*. As a result, (39-a) is reported to be acceptable. In contrast, similar denial of the content of the evidential *allegedly* in (39-b) sounds relatively unacceptable. This represents a third way in which the content encoded in attitude markers seems to differ semantically from canonical not-at-issue content, like that encoded in utterance modifiers, evidentials, and some expressives.

3.2.4 Interim summary

I've delineated class of linguistic markers based on the meaning they encode: the speaker's attitude towards the descriptive content of the utterance. This groups together words like *alas* with prosodic markers like exclamation intonation; it also includes some discourse particles like the Finnish *pä*, and the mirative interpretation of mirative evidentials. And it excludes many other linguistic elements that are prima facie similar: evidentials, utterance modifiers like *frankly*, expressives, appositives, and speaker certainty markers.

Like these other elements, attitude markers encode not-at-issue content: their semantic contribution cannot be targeted by truth-conditional operators, be denied in discourse, or used to address the Question Under Discussion. But I've argued here that attitude markers don't behave like other encoders of not-at-issue content in at least three respects.

First, attitude markers, to the exclusion of the others, are incompatible with non-declarative mood. This is true, with one slight exception, of attitude markers like *unfortunately* in English; of discourse particles like *pä* in Finnish; and, most starkly, of the mirative but not evidential interpretation of the polysemous mirative evidential in languages like Cheyenne. Second, attitude markers take narrower scope than do other not-at-issue encoders. When they occur in the antecedent of a conditional, they must apply to the antecedent, while utterance modifiers and evidentials in the same position instead apply to the conditional as a whole.

Finally, I have extended an empirical claim about mirative evidentials and Moore's Paradox from Murray (2010) to attitude markers more generally. There is some reason to think that denying the content of an attitude marker results in Moore's Paradox, while denying the content of other encoders of not-at-issue content does not. The case is clearest in the comparison of attitude markers and evidentials in (39): embedding an attitude marker and its denial under *suppose* removes the sense of Moorean Paradox, while embedding an evidential and its denial does not.

In what follows, I will propose an account of attitude markers in which they

effectively restrict the sincerity conditions of an utterance, which is also, in part, a role played by illocutionary mood. I couch the analysis in a dynamic update semantics – based on Farkas and Bruce (2010) but drawing heavily from similar adaptations in Murray (2010, 2014); Ettinger and Malamud (2015) – and model illocutionary content as meaning that is added to the speaker’s set of Discourse Commitments (Gunlogson, 2001), rather than the Common Ground (as descriptive not-at-issue content does).

I will argue that this treatment of attitude markers as modeling ‘illocutionary content’ – effectively, as restrictors of sincerity conditions – does a good job of naturally accounting for the semantic differences observed in this section between attitude markers and other encoders of not-at-issue content. It is, additionally, a satisfying model of the persistent intuition – dating at least back to Frege – that attitude markers belong more to the realm of illocution than description.¹¹

I do not, however, attempt to make the stronger claim that no extant theory of canonical not-at-issue content is capable of modeling the idiosyncratic behavior of attitude markers detailed here. Recall that there are two such types of theory: static accounts, (like Potts, 2003a,b, 2005), which treat not-at-issue content in a second, fairly independent semantic tier; and dynamic accounts, (like Murray, 2010, 2014; AnderBois et al., 2010, 2013), which encode not-at-issue content in an automatic common-ground update. These approaches could, in principle, treat attitude markers as they do other encoders of not-at-issue content, and accommodate their differences in behavior a case-by-case basis. But I suspect that such treatments wouldn’t be cohesive, explanatory, or satisfying. In particular, I don’t see how they could predict, as I attempt to do in the next section, a connection between the kind of meaning encoded in attitude markers with their semantic behavior.

4 Modeling illocutionary content

In accordance with the conclusions above, I will outline a proposal here for modeling illocutionary content in a compositional semantics. I use the term “illocutionary content” to refer to not-at-issue meaning, encoded lexically or prosodically in an utterance, pertaining to how the denotation of a sentence affects or is applied to the context. My claim here is that illocutionary content includes (but is not necessarily limited to) the contribution of illocutionary mood and of attitude markers. The claim that attitude markers contribute to illocutionary content is evidenced in part by their discrimination of illocutionary mood, and their ability to result in Moore’s Paradox.

I will argue that attitude markers behave the way they do because they effectively restrict the sincerity conditions of an utterance; and that they restrict

¹¹Frege (1956, 295) says, “An indicative sentence often contains, as well as a thought and the assertion, a third component over which the assertion does not extend. This is often said to act on the feelings, the mood of the hearer or to arouse his imagination. Words like *alas* or *thank God* belong here.”

the sincerity conditions of an utterance because they encode *speaker-oriented propositional attitudes*. Because they mean what they do, attitude markers a) are discourse-anaphoric to a salient proposition associated with the utterance; and b) encode not-at-issue content about the speaker rather than the world. In the present analysis, attitude markers behave differently from other propositional encoders of not-at-issue content because the latter add content to the Common Ground, while attitude markers, instead, add to the speaker’s Discourse Commitments.

The analysis is presented as follows: in §4.1 I will introduce the formal foundations for the analysis, based largely on the theory in Farkas and Bruce (2010); in §4.2 I will adapt Farkas and Bruce’s framework to model what I informally describe as attitude markers’ restriction of the sincerity conditions of an utterance, which accounts for their sensitivity to illocutionary mood type and Moore’s Paradox; in §4.3 I will incorporate focus semantics to account for the idiosyncratic scopal behavior of attitude markers.

4.1 A formal foundation

The claims made here – that sincerity conditions can be restricted by lexical or prosodic markers – underscores the need for a compositional semantic representation of sincerity conditions, or something equivalent. The compositional semantics developed below distinguishes between at-issue content, not-at-issue content, and illocutionary content; formally models sincerity conditions; and represents salience in a way that can include focus alternatives and not-at-issue content. While the proposal is consistent in principle with a number of different formal semantics for illocutionary mood (including Gazdar, 1976; Asher and Lascarides, 2003; Lascarides and Asher, 2009; Krifka, 2001, 2014; Murray, 2014; Murray and Starr, 2016; Portner, 2016), I adopt the framework proposed in Farkas and Bruce (2010) as a foundation, largely because it is flexible enough to easily supplement with my other theoretical desiderata.

Farkas and Bruce (2010) define speech acts as functions from input discourse structures K_i to output discourse structures K_o ; a subcomponent of any discourse structure is a (possibly empty) set of propositions that are under consideration for addition to the CG. This set is called a projection set, and an assertion that p adds p to the input ps . They specify that interrogative mood differs from declarative mood in that the former adds a non-singleton set of propositions to the ps , while the latter adds and projects only one proposition (p.88).

They use the notion of a stack (i.e. a Table T) to model salience in discourse (Ginzburg, 1996; Roberts, 1996). In addition to affecting the common ground, utterances can also raise propositional discourse referents (“drefs”) to salience, and they do so in this theory by pushing the drefs to the top of the Table. Their original conception of the Table separates at-issue content from not-at-issue content; I will revisit this assumption in §4.3.

One other important innovation of Farkas and Bruce’s approach is the incorporation of the notion of a participant’s discourse commitments as a set of

propositions tracked throughout the conversation but distinct from the CG. The innovation is borrowed from Walker (1996) and Gunlogson (2001). They characterize the distinction as “propositions [the speaker] has publicly committed to in the course of [the conversation] and which have not (yet) become mutual commitments” (85). A participant’s ‘total discourse commitments’ is the union of her DC set with the CG.

While Farkas and Bruce define the CG in terms of the discourse commitments of the discourse participants (specifically, as the intersection of all DC sets), I will instead characterize it independently, as the set of propositions that constitute the mutual beliefs of the discourse participants (for the purpose of the conversation).¹²

To summarize, the theory in Farkas and Bruce (2010) relies on characterizing several different subcomponents of a given discourse structure K :

1. the **common ground** (CG), the set of propositions believed by all the discourse participants (for the purpose of the conversation);
2. sets of **discourse commitments** (DC): for each participant x , the set of propositions x has publicly committed to during the conversation;
3. the **Table** T , modeling discourse salience;
4. the **projection set** (ps), the set of beliefs that are being considered for addition into the CG.

Farkas and Bruce (2010) adopt from Krifka (2001) a particular formulation of illocutionary mood in which it takes a sentence as its argument and outputs a function from input to output context states. The declarative mood D is defined over an indicative sentence S_p , a speaker or author a and a discourse structure K_i ; its output is a discourse structure K_o such that K_o is restricted as in (40) (Farkas and Bruce, 2010, 92). (40) has been modified slightly for terminological consistency, and I’ve labeled it “to be revised” because I will amend it on the next page to explicitly differentiate between at-issue and not-at-issue content.

$$\begin{aligned}
 (40) \quad & \text{Declarative operator } (\mathbf{D}), \text{ for sentences } S_p \text{ with at-issue content } p: \\
 & \mathbf{D}(S_p, a, K_i) = K_o \text{ such that} \quad \quad \quad (to \text{ be revised}) \\
 & \quad (i) \quad DC_{a,o} = DC_{a,i} \cup \{p\} \\
 & \quad (ii) \quad T_o = \text{push}(\langle S_p; \{p\} \rangle, T_i) \\
 & \quad (iii) \quad ps_o = ps_i \sqcup \{p\}
 \end{aligned}$$

Step (i) in (40) models the addition of the at-issue content p (the propositional content of the sentence S_p) to the set of propositions representing the speaker’s discourse commitments: those propositions the speaker has publicly

¹²This revision follows criticisms in Šafářová (2007); Poschmann (2008); Gunlogson (2008) that Farkas and Bruce’s CG, as it’s formulated, mistakenly grants one participant the power to commit another participant to a proposition. Thanks to an anonymous reviewer for pointing out the need for a change here.

committed to during the conversation. Step (ii) in (40) represents that the utterance of S_p makes salient the proposition p ; it defines the output stack T_o as the input stack T_i with p pushed on top. Step (iii) represents the illocutionary content of assertion, using the notion of an input (ps_i) and output (ps_o) projection set.¹³ These propositions can then be added to the CG or eliminated throughout the discourse.

Farkas and Bruce’s polar question operator **PQ** takes an interrogative sentence $S_{p?}$ and a discourse structure K_i as its arguments. It raises the issue of whether p by adding p and $\neg p$ to the stack, in (41) (i). It proposes, in (ii), that the interlocutors accept either that p or that $\neg p$ (Farkas and Bruce, 2010, 95).

$$(41) \quad \text{Polar question operator (PQ)} \quad \text{(to be revised)}$$

$$\mathbf{PQ}(S_{p?}, K_i) = K_o \text{ such that}$$

$$(i) \quad T_o = \text{push}(\langle S_{p?}; \{p, \neg p\} \rangle, T_i)$$

$$(ii) \quad ps_o = ps_i \sqcup \{p, \neg p\}$$

The relevant difference between the two illocutionary moods is that declarative mood introduces a singleton set of propositions, while the polar question introduces a non-singleton. These illocutionary moods, as they’re defined, make predictions about the sort of content an utterance makes salient (e.g. a proposition p) and the sort of effect that utterance has on the common ground, by virtue of what it adds to the projection set. This allows Farkas and Bruce to model discourse anaphora to propositions, as many dynamic accounts do, as well as the discourse effect of question responses, which can result, in part, in the acceptance of projection set propositions to the common ground.

The claim that assertions effectively *propose* to update the CG (instead of directly updating it) comes, as far as I can tell, from Clark 1992; Ginzburg 1996 and is motivated by the observation that assertions can be denied in discourse in a way that not-at-issue content cannot.¹⁴ As a result, several recent theories have distinguished between at-issue and not-at-issue content by treating the latter as directly updating the common ground.

Murray (2010, 2014) argued that Cheyenne evidentials introduce not-at-issue content; her semantic theory analyzed not-at-issue content as a direct CG update. (See Portner, 2006; AnderBois et al., 2010, 2013, for similar proposals.) We can supplement Farkas and Bruce’s illocutionary mood operators in (40) and (41) with Murray’s treatment of not-at-issue content by adding a line that a not-at-issue proposition q directly updates the CG.

$$(42) \quad \text{Declarative operator (D)}, \text{ for sentences } S_p \text{ with at-issue content } p \text{ and}$$

$$\text{not-at-issue content } q: \text{ (to be revised)}$$

$$\mathbf{D}(S_p, a, K_i) = K_o \text{ such that}$$

$$(i) \quad DC_{a,o} = DC_{a,i} \cup \{p\}$$

$$(ii) \quad T_o = \text{push}(\langle S_p; \{p\} \rangle, T_i)$$

¹³ \sqcup represents set union minus the elimination of inconsistent propositions (p90).

¹⁴Another motivation is the intuition that, in asserting that p , a speaker cannot force her addressee to commit to p (Malamud and Stephenson, 2014).

- (iii) $ps_o = ps_i \sqcup \{p\}$
 - (iv) $CG_o = CG_i \cup \{q\}$
- (43) *Polar question operator (PQ)*, for an interrogative sentence $S_{p?}$ with at-issue content p , $\neg p$ and not-at-issue content q :
PQ($S_{p?}, K_i$) = K_o such that
- (i) $T_o = \text{push}(\langle S_{p?}; \{p, \neg p\} \rangle, T_i)$
 - (ii) $ps_o = ps_i \sqcup \{p, \neg p\}$
 - (iii) $CG_o = CG_i \cup \{q\}$

In sum, I’ve adopted the semantic treatment of illocutionary mood in Farkas and Bruce (2010) to form the foundation of the account. It encodes illocutionary mood – at least for declaratives and polar questions – and additionally models salience and discourse commitments, which I’ll take advantage of for the formal proposal in §4.2. But while Farkas and Bruce characterize at-issue content as proposals to update the common ground, they do not include an explicit treatment of not-at-issue content. I’ve adapted their account, following Murray (2010), to treat not-at-issue content as direct common ground update.

This covers the first two desiderata outlined above: we have in hand a theory that represents illocutionary mood and distinguishes between at-issue and not-at-issue content. In what follows, I’ll supplement this theory further.

4.2 Sincerity conditions in an update semantics

Farkas and Bruce (2010) foresee the need to expand their analysis:

“For the matters we discuss here, further additions to context structure such as the agendas of participants or representations of their private doxastic states are not necessary. The model we provide is consistent with expansion in these directions, as well as with additions of finer-grained structures for dealing with anaphoric relations.” (Farkas and Bruce, 2010, 89)

The goal of this subsection is to expand Farkas and Bruce’s framework in just this way. I’ll begin by discussing how we can use their formalism to model speakers’ attitudes; I’ll expand the notion of Discourse Commitments to account for the sorts of content attitude markers can apply to.

According to Searle (1969), the utterance of an assertion is conventionally associated with several constitutive rules, one of which is the sincerity condition that the speaker believes the content of the utterance. An utterance that does not satisfy this condition is uttered insincerely; this is one way in which sincerity conditions differ from other constitutive rules, like preparatory conditions. And in fact, it seems appropriate to characterize the utterance of a sentence containing an attitude marker (e.g. *alas*) by a speaker who is not in fact dismayed by its content as an insincere utterance.

We have already seen one other reason to associate the content of attitude markers with sincerity conditions: their participation in Moore’s Paradox. Sincerity conditions have been identified by some, including Searle (1969, 65), as

the culprit behind Moore’s Paradox. In particular, an appealing explanation of Moore’s Paradox is that it results when the speaker denies the content of a sincerity condition on her utterance. If attitude markers contribute sincerity conditions – like illocutionary mood does – we would expect the denial of their content, too, to result in Moore’s Paradox.

Following Searle (1969) (see especially his discussion of promising, p.57), I will assume that this additional sincerity condition effects changes in the essential conditions (i.e. the discourse properties) of the utterance. The essential conditions of an exclamation, then, effectively include “expresses that *S* is surprised that *p*”.

Searle and other Speech Act theorists have considered sincerity and essential conditions to be properties of a speech act. However, recent dynamic proposals follow Stalnaker (1973, 1978) in characterizing the effect of an utterance on the discourse as a property of illocutionary mood instead (Murray, 2014; Murray and Starr, 2016). Such an approach offers a compositional account of illocutionary mood in a way that still allows it to not fully determine illocutionary force. The current proposal is that we should similarly adapt Speech Act Theorists’ treatment of sincerity conditions: instead of being a property of a speech act, sincerity conditions are a part of the sentence’s content, introduced by illocutionary mood and attitude markers. If this is right, illocutionary mood and attitude markers – qua illocutionary content – have in common that they restrict a speaker’s Discourse Commitments.

This perspective has the benefit of treating illocutionary mood and attitude markers compositionally, while still allowing for a many-to-many relationship between illocutionary mood and illocutionary force (see, among others, Harnisch, 2005). In particular, from this perspective, an utterance of the declarative sentence *Your behavior has, alas, made me angry* can count as a threat when contextual information about the context of utterance supplements the illocutionary content it encodes semantically: namely, that it updates the common ground with the proposition ‘Your behavior has made me angry’ (a contribution of its illocutionary mood); and that its use is licensed when, and thereby signifies that, the speaker is dismayed by that proposition (the contribution of the attitude marker).

Informally, my analysis characterizes attitude markers as restricting the sincerity conditions of an utterance with the requirement that the speaker be surprised or dismayed (depending on the attitude marker) by that sentence’s descriptive content. The connection between attitude modifiers and sincerity conditions is explicit in Vanderveken (1990), according to whom *alas* expresses “the sincerity condition that the speaker is unhappy with the existence of that state of affairs” (p128). However, while Vanderveken (1990) treated attitude markers as modifiers of illocutionary mood (restricting the sincerity conditions of that mood), I analyze them as restricting the speaker’s Discourse Commitments independently of illocutionary mood.

The original characterization of Discourse Commitments in Gunlogson (2001) is in terms of beliefs that each participant is publicly committed to: “public in the sense that the participant is mutually recognized as committed to them”

(p42). She defines them, crucially, in terms of beliefs:

- (44) **Discourse commitments** *(to be revised)*
Let DC_a and DC_b be sets of propositions representing the public beliefs of a and b , respectively, with respect to a discourse in which a and b are the participants, where:
- a. p is a public belief of a iff ‘ a believes p ’ is a mutual belief of a and b
 - b. p is a public belief of b iff ‘ b believes p ’ is a mutual belief of a and b

While they were initially proposed to address speaker bias in rising declaratives, public commitments are additionally useful for modelling sincerity conditions. Historically, speech act theorists have differed with respect to whether they encode illocutionary force in speaker (and hearer) commitments (e.g. Krifka, 2014), or whether they encode it in terms of what Asher and Lascarides (2008) call ‘BDI’: beliefs, desires, and intentions. In treating sincerity and essential conditions as part of an utterance’s “illocutionary content,” encoded semantically, I follow Gunlogson (2001) and others’ proposal to adopt public commitments instead of Searle’s constitutive rules.

Harnish (2005) reviews a number of objections to *reducing* illocutionary force to commitments – as Krifka (2014) does in his approach (in which speech acts are typed as “commitment change potentials”) – but the division of labor proposed here side-steps most of these problems. The idea is that what Searle thought of as the speaker’s belief in p is in fact part of the speaker’s discourse commitments: in addition to its effect on the common ground, an act of assertion that p publicly commits the speaker to p . There are, technically speaking, differences between Searle’s sincerity condition on assertion (that the speaker believe that p) and the Gunlogson/Farkas/Bruce characterization of an assertion adding p to the speaker’s set of discourse commitments. This is because publicly committing to a proposition p doesn’t reduce to believing that p (and vice versa).

However, a speaker’s belief that p and her public commitment to p are closely enough related that we can treat public commitment as a proxy for belief – for the purposes of modeling conversation – and additionally use DC s to encode the contribution of attitude markers. Specifically, a speaker’s discourse commitments are things that the speaker is committed to treating, for the purposes of the conversation, as if she did believe them. In other words: in contexts in which the speaker is being sincere (or in which the hearer assumes the speaker is sincere), the speaker’s publicly committing to p amounts to the speaker’s assurance she believes that p . If this is right, then the assertion operator proposed by Farkas and Bruce (2010) in (42) represents, albeit indirectly, Searle’s sincerity condition on assertion.¹⁵

¹⁵The equivalence is not a perfect one: it’s easy to imagine a situation in which each discourse participant believes that p , and knows the others believe that p , but in which no one wants to publicly commit to p : say, a situation in which the boss has some toilet paper stuck to her shoe. Of course, it’s possible that Discourse Commitments model the sincerity conditions on assertion more accurately than Searle’s does, in which case the differences between public commitment and belief are unproblematic. Arguments to this effect can be found in Asher

Recall that Gunlogson defines discourse commitments in terms of beliefs: “ p is a public belief of a iff ‘ a believes p ’ is a mutual belief of a and a ” (p42). If attitude markers restrict the sincerity conditions of an utterance, and if attitude markers encode propositional attitudes in addition to belief, we will need to modify this notion of discourse commitments to include other relations between a speaker and a proposition. This definition can be extended to DC_b , as Gunlogson’s is.¹⁶

- (45) **Discourse Commitments** (final)
 Let DC_a be a set of pairs representing the public commitments of a with respect to a discourse in which a and b are the participants, where:
- a. $\langle \text{believes}, p \rangle$ is a public commitment of a iff ‘ a believes p ’ is a mutual belief of a and b ;
 - b. $\langle \text{is-disappointed}, p \rangle$ is a public commitment of a iff ‘ a is disappointed that p ’ is a mutual belief of a and b ; and
 - c. $\langle \text{is-surprised}, p \rangle$ is a public commitment of a iff ‘ a is surprised that p ’ is a mutual belief of a and b .

This switch from a set of propositions in terms of speaker belief to a set of pairs is reminiscent of a similar innovation proposed in Portner (2006). Portner proposes an account of speaker certainty markers (as encoded in some evidentials and “utterance modifiers” like *honestly*) wherein the set of propositions representing the Common Ground (his *ps*) is subdivided into those mutually agreed to be true (the traditional Common Ground) and those reported to be true, conjectured to be true, etc. Importantly, this proposal introduces no CG subsets; it just alters the CG from a set of propositions to a set of pairs of propositional attitudes and propositions.¹⁷

To track this change in the characterization of Discourse Commitments, I amend the formulation of the sincerity conditions encoded in declarative mood to include a pair (from (42)):

- (46) *Declarative operator* (**D**), for sentences S_p with at-issue content p and not-at-issue content q : (final)
 $\mathbf{D}(S_p, a, K_i) = K_o$ such that

and Lascarides (2003, 2008); Lascarides and Asher (2009), whose specific approach I do not adopt because, like Krifka (2014), they do not distinguish between force and mood.

¹⁶I notate the contents of DC sets as ordered pairs (instead of as, for instance, the apparently equivalent *believes(p)*) because I believe that this notation better reflects the nature of DC sets. In particular, it is a natural way to require that all additions to the DC set contain a propositional attitude and a proposition.

¹⁷There are additional proposals to modify update-semantic theories like Farkas and Bruce’s in order to track other components of discourse. In their analysis of the Mandarin discourse particle *ba*, Ettinger and Malamud (2015) model a request for hearer involvement by introducing different sub-parts of Table to designate different levels of speaker commitment to the proposed CG update. And in his recent theory of the semantics of imperatives and modal particles, Portner (2016) supplements this sort of account with a list of priorities. He says (p14): “just as we must maintain both the common ground and individual commitment slates in our discourse model, we also must keep track of the shared to-do list function and individual participants’ understanding of what priorities each participant is committed to.”

- (i) $DC_{a,o} = DC_{a,i} \cup \langle \text{believes}, p \rangle$
- (ii) $T_o = \text{push}(\langle S_p; \{p\} \rangle, T_i)$
- (iii) $ps_o = ps_i \sqcup \{p\}$
- (iv) $CG_o = CG_i \cup \{q\}$

The final component of the analysis is the treatment of attitude markers themselves. I'll model this account on *alas*, but intend it to be generalizable to, at least, exclamation intonation, Finnish *-pää*, and the mirative interpretation of the Cheyenne mirative evidential.

I define *alas* in the same terms as illocutionary mood in Farkas and Bruce (2010), namely as an operator which takes a clause as its argument and outputs a function from input to output context states.

- (47) *Alas* (**A**), for clauses C with content p : (to be revised)
A(S, a, K_i) = K_o such that
- (i) $DC_{a,o} = DC_{a,i} \cup \langle \text{is-disappointed}, p \rangle$
 - (ii) $T_o = \text{push}(\langle S; \{p\} \rangle, T_i)$

There are several ways in which the definition of attitude markers exemplified in (47) differs from that of declarative mood in (46). First, unlike declarative mood, *alas* does not update the projection set ps , which is how this approach models the assertoric component of declarative mood. This is how attitude markers differ importantly from illocutionary mood. Second, *alas* in (47) does not update the Common Ground, which is how this approach models not-at-issue content. Third, (47) analyzes *alas* as restricting the speaker's DC set with an ordered pair whose first member is the propositional attitude *is-disappointed*, not *believe*. These last two characteristics represent how attitude markers differ importantly from canonical, descriptive not-at-issue content: they update the speaker's Discourse Commitments, rather than the Common Ground. (47) additionally differs in that the proposition p the attitude marker is anaphoric to is less specified than it is in (46) for declarative mood. The nature of this propositional anaphora will be further clarified in the next subsection.

The extent to which Discourse Commitments, in this approach, involve propositional attitudes other than belief is constrained lexically. In this paper, I discuss attitude markers that encode disappointment (e.g. *alas*) and surprise (e.g. mirativity markers); if there is evidence of other attitude markers that encode additional propositional attitudes, they would also be tracked in these Discourse Commitment pairs.

Because, according to (47), attitude markers update the speaker's *DC*, rather than the *CG*, the meaning they encode is importantly distinct from e.g. not-at-issue content. It is of course reasonable to assume that information encoded in Discourse Commitments might or will, at some point, enter the Common Ground and thereby support presuppositions in later discourse. I will leave such a mechanism undefined for now.

Encoding the content of an attitude marker instead in a speaker's Discourse Commitments gives us a way of preserving its not-at-issue status while address-

ing its difference in meaning from that of canonical not-at-issue encoders. I've argued that DC sets are appropriate for modeling the sincerity conditions of an utterance because they represent the speaker's public commitments, and to be insincere is to falsely commit oneself to something publicly. Searle (1969) and others have argued that the unacceptability of Moorean sentences is best attributed to the conflict of one claim with the sincerity conditions of another; in this framework, according to that perspective, Moore's Paradox results when one claim contradicts that speaker's Discourse Commitments. Evidently, some speakers have intuitions about the difference between contradicting information in the Common Ground and contradicting information in a speaker's DC set; since attitude markers behave like classical Moore's Paradox cases in this respect, their content should receive the same formal treatment as (other) sincerity conditions.

This characterization of attitude markers accounts for another aspect of their idiosyncratic semantic behavior: their incompatibility with non-declarative mood. (47) is defined only over proposition-denoting clauses; while the descriptive content of a declarative sentence is a proposition p , the descriptive content of a question is a set of propositions (Hamblin, 1971) (and, arguably, an imperative denotes a property (Hausser, 1980; Portner, 2004; Murray and Starr, 2016) or some other non-propositional content). This definition predicts, correctly, that attitude markers can't apply to a (matrix) question or imperative. It also correctly predicts that (lexical) attitude markers are acceptable in clauses embedded in questions, as (48) shows.

(48) What does Sue, who alas couldn't be here today, think about the proposal?

The next subsection discusses how the proposal in (47) can be extended to predict that attitude markers necessarily take narrow scope in these embedded contexts.

There's a reasonable perspective from which this isn't stipulative: it's not clear that it's possible to hold a propositional attitude towards a set of propositions. As a reviewer points out, even explicitly encoded speaker attitudes are prohibited in questions, presumably for this reason:

(49) *Has John arrived on time, which disappointed me?

In contrast, attitude markers are acceptable in tag questions, which are typically associated with propositional content (as well as speaker bias, Ladd, 1981; Romero and Han, 2004; Gunlogson, 2001; Malamud and Stephenson, 2014).¹⁸

(50) Alas, John has arrived on time, hasn't he?

As a result, at least given the characterization of the polar question operator **PQ** in (41), the analysis of attitude markers in (57) predicts that they are unacceptable in questions, or any other construction associated with non-propositional

¹⁸Thanks to Amy Rose Deal (p.c.) for pointing out the significance of these data.

content.¹⁹

There is additional evidence that this is a fortunate result. Recent work in alternative semantics (Alonso-Ovalle, 2006, among others) and inquisitive semantics (Groenendijk, 2009, among others) have proposed treating certain cases of disjunction as similar to polar questions in just this respect. Alonso-Ovalle (2006), for instance, adopts a Hamblin semantics in which, in certain contexts, disjunctive sentences denote a set of propositions (i.e. multiple alternatives). If these approaches are right, (47) predicts that e.g. *alas* is unacceptable in alternative-projecting disjunctive sentences in just the same way they’re unacceptable in polar questions. And this seems to be the case:

- (51) a. #(Wow,) John rode his bike or arrived on time!
b. #Alas, John rode his bike or arrived on time.

This is not to say that attitude markers are ungrammatical in a sentence that includes a disjunction, just in those in which the disjunction introduces multiple alternatives. It’s possible, of course, to be disappointed at a single, disjunctive proposition; in a situation in which you learn that John’s family is no longer a two-income family, it is perfectly acceptable to lament *Alas, John lost his job or Mary lost hers*.²⁰ These sentences differ intonationally in English, and involve distinct lexical items in languages like Egyptian Arabic (Winans to appear).

However, recall that §3.2 suggested a possible difference within the class of attitude markers:

- (52) a. *Unfortunately, did John lose the race?
b. Alas, did John lose the race?

I reported the intuition that, in (52-b), *alas* is anaphoric on the bias of the question (that John lost the race), rather than the content of the question (the set of propositions {John lost the race, John didn’t lose the race}). While accounts that differentiate between these two things semantically (Romero and Han, 2004; Gunlogson, 2001; Malamud and Stephenson, 2014) could explain the data in (52-b) even given the analysis in (57), I currently have no explanation for how two attitude markers could differ in their ability to be anaphoric on a question’s bias. Perhaps, as suggested in §3.2, the difference is best attributed to diachronic differences between the two attitude markers.

The paper began with the claim that attitude markers form a natural class within the larger group of encoders of not-at-issue content by virtue of the meaning they encode: in particular, that they behave differently from other encoders of not-at-issue content because they mean what they mean. I can now partially substantiate that claim: attitude markers behave the way they

¹⁹I am unaware the existence of any question-embedding attitude markers, those lexically encoding attitudes like wonder. Nothing I say here predicts their existence or absence in natural language; however, because they don’t encode propositional attitudes, they wouldn’t be strictly speaking attitude markers according to my definition. I would predict that they, too, would discriminate illocutionary mood, insofar as they would be anaphoric to a set of propositions rather than a single proposition.

²⁰Thanks to an anonymous reviewer for this example.

do because they encode not-at-issue information about the speaker’s propositional attitudes. Because the information is about the speaker’s *propositional* attitudes, it is anaphoric on a (single) proposition, and so is incompatible with illocutionary mood associated with more than one proposition, or not associated with any propositions. And because the information is about *the speaker’s* propositional attitudes is represented as his or her public commitment, instead of being automatically introduced into the Common Ground, as evidenced by the ability of attitude markers to participate in Moore’s Paradoxes. This is, arguably, an intuitive way of modeling the apparent differences between descriptive and illocutionary content (and, possibly, Kaplan’s related distinction between descriptive and expressive content).

In the next section, I argue that attitude markers apply to focus-semantic content of at-issue propositions. This will allow for an account of attitude markers’ focus sensitivity, and also provides some explanation for their unique scopal restrictions.

4.3 Illocutionary mood and focus semantics

The goal of this section is to add one final layer of detail to the analysis in (47) in order to address the nature of propositions that attitude markers can be anaphoric to: in particular, that they are focus-sensitive and must take local or narrow scope when embedded.

As discussed in §3.2 (repeated in (53)), attitude markers are focus-sensitive, which means that they apply to (in this account, are anaphoric to) different propositions, depending on the location of focus.

- (53) a. Alas, JOHN_F got a dog.
 b. Alas, John got a DOG_F.

I account for this by proposing that attitude markers are anaphoric to the ordinary value of an assertion, which includes the negation of focus alternatives. This allows for an account of the discourse-sensitivity of attitude markers.

Rooth (1985) introduced focus-semantic values to allow focus-sensitive operators to explicitly manipulate a proposition’s alternatives. If the ordinary meaning of an element α (notated $\llbracket \alpha \rrbracket^o$) is type $\langle \sigma \rangle$, the focus meaning of that element (notated $\llbracket \alpha \rrbracket^f$) is a set of alternatives to α , type $\langle \sigma, t \rangle$. The alternative set is context-sensitive; the set of focus alternatives C to a proper name like *Mary* could be the set {tom, dick, harry} or it could be the entire domain of individuals.

A focus-sensitive operator like *only* (Rooth, 1985) or a synonymous null exhaustivity operator (Groenendijk and Stokhof, 1984) applies to a sentence’s focus-semantic value – a set of propositions – to produce its ordinary meaning: a proposition that explicitly negates (or eliminates as possibilities) the focus alternatives in C . (55) shows an example of how an ordinary semantic value is computed in this way.

$$(54) \quad \llbracket \text{EXH} \rrbracket = \llbracket \text{only} \rrbracket = \lambda C \lambda p \forall q [(q \in C \wedge \text{true}(q)) \leftrightarrow q = p]$$

- (55) a. $\llbracket \text{Mary}_F \text{ likes Sue} \rrbracket^f = \{\text{likes}(t,s), \text{likes}(d,s), \text{likes}(h,s)\}$
 b. $\llbracket \text{EXH Mary}_F \text{ likes Sue} \rrbracket^f =$
 $\forall q[q \in \{\text{likes}(t,s), \text{likes}(d,s), \text{likes}(h,s)\} \wedge \text{true}(q) \leftrightarrow q = \text{likes}(m,s)]$

I am not the first to associate attitude markers with focus-semantics. Castroviejo-Miró (2008) proposes an alternative-based semantics for *wh*-exclamatives in which they denote a set of ordered alternatives, and Anderson (2016) proposes an alternative-based account of English *some*-exclamations (e.g. *Some friend she turned out to be!*) in which the alternatives are triggered by the epistemic indefinite *some* (Kratzer and Shimoyama, 2002; Alonso-Ovalle and Menéndez-Benito, 2003). His account additionally involves a syntactically represented exclamation intonation marker ‘ExOp’ that structures the alternatives and “asserts an attitude towards the most unexpected proposition”. The present proposal makes more explicit the semantics behind exclamation intonation and other attitude markers, but has in common with these accounts that attitude markers is sensitive to the set of focus alternatives of an utterance. Crucially, however, I do not propose that attitude markers operate on sets of propositions, like Castroviejo and Anderson do; I propose that they apply to a proposition that incorporates the focus-semantic value of a sentence, namely the ordinary value of that sentence. Importantly, the ordinary value of a sentence includes the negation of the relevant alternatives.²¹

In the analysis of attitude markers in §4.2, repeated below, *alas* is anaphoric to some unspecified proposition *p*. In the revised version in (57), I specify that *p* applies to the ordinary value of the clause the attitude marker applies to, which, again, includes the negation of the focus alternatives of the clause.²²

- (57) *Alas* (**A**), for clauses *C* with the ordinary focus-semantic value *p*: (*final*)
 $\mathbf{A}(S, a, K_i) = K_o$ such that
 (i) $DC_{a,o} = DC_{a,i} \cup \langle \text{is-disappointed}, p \rangle$
 (ii) $T_o = \text{push}(\langle S; \{p\} \rangle, T_i)$

Like (47), (57) analyzes *alas* as adding to the speaker’s discourse commitments the proposition that the speaker is disappointed in *p*. The difference is that (57) explicitly represents *p* as the ordinary meaning of the sentence modified by the attitude marker. (58) gives an example of this analysis applied to the example in (1).

²¹(56) shows how the ordinary meanings of the propositions in (53) differ in a sample context of evaluation *c*; in this context, the two utterances would have the same semantic schema as in (58), but they would differ in this way in the value of their argument *p*.

- (56) a. $\llbracket \text{JOHN}_F \text{ got a dog} \rrbracket^{o,c} = \exists x[\text{dog}(x) \wedge \text{got}(\text{john},x)] \wedge \neg \exists x[\text{dog}(x) \wedge \text{got}(\text{mary},x)] \wedge$
 $\exists x[\text{dog}(x) \wedge \text{got}(\text{sue},x)]$
 b. $\llbracket \text{John got a DOG}_F \rrbracket^{o,c} = \exists x[\text{dog}(x) \wedge \text{got}(\text{john},x)] \wedge \neg \exists x[\text{cat}(x) \wedge \text{got}(\text{john},x)] \wedge$
 $\neg \exists x[\text{chinchilla}(x) \wedge \text{got}(\text{john},x)]$

²²A more sophisticated presentation of the same requirement could explicitly reference a null exhaustivity operator, either as a precondition on the application of the attitude marker, or in terms of some selectional restriction. I encode the requirement in the metasemantics because I am neutral about its particular implementation.

- (58) $\llbracket \text{Alas, John lost the race} \rrbracket = \mathbf{D}(\mathbf{A}(S, a, w_{\text{@}}, K_i)) = K_o$ such that
- (i) $DC_{a,o} = DC_{a,i} \cup \langle \text{believes}, p \rangle$
 - (ii) $T_o = \text{push}(\langle S; p \rangle, T_i)$
 - (iii) $ps_o = ps_i \sqcup \{p\}$
 - (iv) $CG_o = CG_i \cup \{q\}$
 - (v) $DC_{a,o} = DC_{a,i} \cup \langle \text{is-disappointed}, p \rangle$

The first four lines are the result of the application of the declarative mood marker \mathbf{D} ; the final line is the contribution of *alas*. It adds to the speaker's Discourse Commitment set the proposition that the speaker is disappointed that p which, in this simple example, is equivalent to the at-issue content of the utterance (i.e. the proposition that John lost the race).

In later work (1992; 1996), Rooth connected the role of focus-semantic values to discourse coherence and, in particular, question-answer congruence. Roberts (1996) expanded this idea into a larger theory of discourse coherence based on QUDs (see also Carlson, 1983). In her approach, a conversational move β is congruent to a question $?a$ if the set of focal alternatives to β is identical to the set of Q-alternatives that constitutes the semantics of $?a$ (p24). This definition drives her account of prosodic focus: a focused element β carries the presupposition that it is a congruent answer to the relevant QUD.

Given this connection, the sensitivity of attitude markers to focus entails that attitude markers are discourse-sensitive. This is reflected in the definition in (57), which specifies that attitude markers apply to the same proposition that is added to the Table in the dynamic update of the clause they occur in. In other words, defining the scope of an attitude marker in terms of the ordinary focus-semantic value of the clause it occurs in assures that the attitude marker is discourse-sensitive, and thereby operates on the same content that is added to the Table in the relevant dynamic update.

As discussed in §3.2, attitude markers can range over not-at-issue content as well as at-issue content (59).

- (59) John, who *alas* won the lottery, lost the race.

This is part of a larger pattern: when attitude markers occur in an embedded clause, they scope over the embedded clause, rather than the matrix sentence. This is demonstrated by the contrast between (60-a) and (60-b) below (from (25) in §3.2).

- (60) a. *Alas*, if the mayor is convicted, she must resign from office.
 b. If, *alas*, the mayor is convicted, she must resign from office.

The definition in (57) accounts for this restriction because it ties the attitude marker to the focus-semantic ordinary value of the clause it occurs in... which is to say, the most salient proposition (the one added to the Table) in the dynamic update the attitude marker participates in. Attitude markers contribute their own restriction to the speaker's DC set, which (in the case of embedded clauses) requires a sub-sentential dynamic update. As a result, they are

discourse-anaphoric to the proposition denoted by the clause they occur in, not the proposition denoted by the matrix sentence.

However, in contrast to the other properties of attitude markers discussed here, nothing in this account predicts that attitude markers *must* take local scope, by virtue of their meaning. In other words, the analysis in (57) predicts that attitude markers take local scope in embedded clauses because the denotation of the embedded clause is likely to be the most salient proposition in that update. It allows for (but does not account for) at least two types of deviation from the characterization of the scope of attitude markers made in Section 3.2: attitude markers that can be embedded but, in certain contexts, need not apply to the embedded clause; and attitude markers that cannot be embedded. I'll end by discussing these predictions in turn.

Recall from §3.2 that *alas* differs from other attitude markers in its ability to occur in polar questions (but not constituent questions):

- (61) a. *Unfortunately, did John lose the race?
b. Alas, did John lose the race?

Intuitively, the attitude marker in (61-b) is not applying to the question itself (nor can it, according to the account here); it is applying to one of the two propositions denoted by the question; in particular, the one to which the speaker is biased, according to work in Romero and Han (2004); Reese (2008), i.e. the proposition that John lost the race.

I cannot account for this difference between *unfortunately* and *alas*²³, but this lexical variation is allowed by the current proposal, which specifies that the scope of an attitude marker is the focus-semantic ordinary value of the clause it occurs in. This definition allows, in principle, discourse anaphora to the biased proposition in the denotation of a question, in an account of bias that makes it salient (enough).

Recall also that the Pottisian expressive *damn*, which classifies as an attitude marker when it ranges over propositions, cannot be embedded in conditionals at all:

- (62) *If, damn, the mayor is convicted, she'll have to resign!

This, too, is consistent with the above proposal, although again I do not account here for the lexical variation. For whatever reason, there is some lexical quirk of the expressive *damn* that prevents it from occurring in an embedded clause when it applies to a proposition.²⁴ But we have no evidence of an attitude marker

²³Although it's been suggested to me that lexical differences between attitude markers are best attributed to diachronic factors; in particular, that *alas* is more distantly related to its adjectival or adverbial roots than *unfortunately* (e.g. *It was unfortunate/*alas that John lost the race*).

²⁴I will, however, observe that unlike *alas* and the other attitude markers discussed here, *damn* can occur on its own, without any overt propositional argument. Its possible, then, that this means *Damn, John lost the race!* is better notated as *Damn! John lost the race!* in which case it might differ slightly from the otherwise relatively uniform class of attitude markers discussed here.

which can be embedded scoping over anything other than the proposition added to the Table in the relevant dynamic update.

These observations raise the issue of how to best analyze encoders of descriptive not-at-issue content, like the evidential adverb *apparently*, which necessarily scope outside of their embedded clause. It’s likely, given the discussion here, that these operators are less discourse-sensitive than are attitude markers, although it’s unclear to me how to best model this difference.

5 Conclusions

In this paper, I’ve used attitude markers as a case study for better examining the traditional divide between descriptive and illocutionary (or expressive, in the sense of Kaplan, 1997) content. I’ve argued that it differs from the recent distinction between at-issue and not-at-issue content. The meaning encoded in attitude markers, while clearly not-at-issue, differs in several ways from that of canonical encoders of not-at-issue content, like utterance modifiers, evidential adverbs, appositives, etc.: it is incompatible with non-declarative mood; its scope is restricted in certain contexts; and it can participate in Moore’s Paradox.

I’ve presented a formal account of attitude markers in which the content they apply to is added to the speaker’s Discourse Commitments in the form of an ordered pair, e.g. $\langle \text{disappointed}, p \rangle$. In this account, attitude markers differ from canonical encoders of not-at-issue content (here, encoders of ‘descriptive not-at-issue content’) in that the information they encode updates the DC set, instead of the Common Ground. And while illocutionary mood restricts the DC set as well – encoding, in the case of declarative mood, the sincerity condition that the speaker believe the content of the utterance – attitude markers differ from mood in this analysis because mood also alters the projection set. I’ve argued assimilating the content of attitude markers with the sincerity conditions of illocutionary mood correctly predicts their ability to participate in Moore’s Paradox; it also correctly predicts that attitude markers can only apply to clauses that denote single propositions.

I’ve defined the class of attitude markers based on the sort of meaning they encode, but I’ve demonstrated that they behave as a natural class semantically, as well, and that this behavior holds of attitude markers across languages. As a result, I’ve argued that attitude markers behave the way they do *because of* the sort of meaning they encode. Because they encode propositional attitudes, they must range over (single) propositions, and in particular, they must range over the most salient proposition encoded in the utterance. I’ve analyzed this as a particular sort of focus-sensitivity, and I’ve argued that it explains attitude markers’ restriction to declarative sentences (or sentences associated with only one salient proposition) as well as its unique, discourse-sensitive scopal restrictions. And because attitude markers encode *the speaker’s* propositional attitude, they add information to the speaker’s Discourse Commitments, like sincerity conditions do, rather than the Common Ground.

To make these arguments, I’ve drawn in part from the phenomenon of mi-

rative evidentials: morphemes that have an evidential interpretation in some contexts, and a mirative interpretation in others (in which case they count as attitude markers). When Cheyenne mirative evidentials act as evidentials, they can occur with non-declarative mood (Murray, 2010; Rett and Murray, 2013). When they act as miratives, however, they cannot. There is currently a debate in the evidentials literature about how to classify evidentials, and this debate is complicated by observations that the compatibility of evidentials with non-declarative mood appears to vary cross-linguistically; in languages like Abkhaz, Baniwa, and Jarawara, evidentials, too, are unacceptable in questions (Aikhenvald, 2004).

While there are lots of open questions about what evidentials are and what they contribute to a sentence, the discussion here offers one possibility of reconciling the observed cross-linguistic differences with respect to compatibility with illocutionary mood. It's possible that, while evidentials in some languages are classified as encoding descriptive not-at-issue content, and thereby update the Common Ground (as they do in Cheyenne; Murray, 2010), evidentials in other languages are classified as encoding illocutionary content about speakers' doxastic states, and thereby update the DC set (as they do, with some important differences, in the account of Cuzco Quechua evidentials in Faller, 2002). In other words, while it's clear that certain aspects of meaning (like the speaker's attitude towards the descriptive content of an utterance) is best encoded as illocutionary content, languages might differ on how other aspects of meaning (like evidence type, or speaker certainty) should be encoded, suggesting that the class of illocutionary content might be to some extent cross-linguistically variable.

References

- Aikhenvald, A. (2004). *Evidentiality*. Oxford University Press.
- Alonso-Ovalle, L. (2006). *Disjunction in alternative semantics*. PhD Thesis, University of Massachusetts, Amherst.
- Alonso-Ovalle, L. and Menéndez-Benito, P. (2003). Some epistemic indefinites. In *Proceedings of the Northeast Linguistics Society (NELS)*, volume 33.
- AnderBois, S., Brasoveanu, A., and Henderson, R. (2010). Crossing the appositive/at-issue meaning boundary. In Li, N. and Lutz, D., editors, *Proceedings of SALT 20*, pages 328–346. CLC.
- AnderBois, S., Brasoveanu, A., and Henderson, R. (2013). At-issue proposals and appositive impositions in discourse. *Journal of Semantics*. doi: 10.1093/jos/fft014.
- Anderson, C. (2016). An alternatives based account of *some*-exclamatives. LSA talk, January 9.
- Asher, N. and Lascarides, A. (2003). *Logics of conversation*. Cambridge University.
- Asher, N. and Lascarides, A. (2008). Commitments, beliefs and intentions in dialogue.
- Bach, K. (1999). The myth of conventional implicature. *Linguistics and Philosophy*, 22:367–421.
- Bartels, C. (1999). *The intonation of English statements and questions: A compositional interpretation*. Garland.

- Büring, D. and Gunlogson, C. (2000). Aren't positive and negative polar questions the same?
- Carlson, L. (1983). *Dialogue Games: An Approach to Discourse Analysis*. Reidel, Dordrecht.
- Castroviejo-Miró, E. (2008). An expressive answer: Some considerations on the semantics and pragmatics of wh-exclamatives. In *Proceedings from the annual meeting of the Chicago Linguistic Society*, volume 44, pages 3–17.
- Cinque, G. (1999). *Adverbs and functional heads: a cross-linguistic perspective*. Oxford University Press.
- Clark, H. (1992). *Arenas of language use*. University of Chicago.
- Cruttenden, A. (1986). *Intonation*. Cambridge University Press.
- DeLancey, S. (1997). Mirativity: the grammatical marking of unexpected information. *Linguistic Typology*, 1:33–52.
- DeLancey, S. (2001). The mirative and evidentiality. *Journal of Pragmatics*, 33:369–382.
- Dickinson, C. (2000). Mirativity in Tsafiki. *Studies in Language*, 24:379–422.
- Ettinger, A. and Malamud, S. (2015). Mandarin utterance-final particle ba in the conversational scoreboard. In Csipak, E. and Zeijlstra, H., editors, *Proceedings of Sinn und Bedeutung 19*.
- Faller, M. (2002). *Semantics and Pragmatics of Evidentials in Cuzco Quechua*. PhD Thesis, Stanford University.
- Faller, M. (2014). Reportativity, (not-)at-issueness, and assertion. In *Proceedings of the fortieth annual meeting of the Berkeley Linguistics Society*, pages 62–84.
- Farkas, D. and Bruce, K. (2010). On reacting to assertions and polar questions. *Journal of Semantics*, 27:81–118.
- Frege, G. (1956). The thought: a logical inquiry. *Mind*, 65:289–311.
- Gazdar, G. (1976). On performative sentences. *Semantikos*, 1/3:37–62.
- Ginzburg, J. (1996). Dynamics and the semantics of dialogue. In Seligman, J. and Westerståhl, D., editors, *Language, Logic and Computation*, volume 1, pages 221–237. Stanford University.
- Groenendijk, J. (2009). Inquisitive semantics: two possibilities for disjunction. In Bosch, P., Gabelaia, D., and Lang, J., editors, *Proceedings of the 7th International Tbilisi Symposium on Language, Logic and Computation*, pages 80–94. Springer-Verlag.
- Groenendijk, J. and Stokhof, M. (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, University of Amsterdam.
- Gunlogson, C. (2001). *True to Form: Rising and Falling Declaratives as Questions in English*. Routledge.
- Gunlogson, C. (2008). A question of commitment. *Belgian Journal of Linguistics*, 22:101–136.
- Gutzmann, D. (2015). *Use-conditional meaning*. Oxford University Press.
- Hamblin, C. (1971). Mathematical models of dialogue. *Theoria*, 37:130–155.
- Harnish, R. (2005). Commitments and speech acts. *Philosophica*, 75:11–41.
- Hausser, R. (1980). Surface compositionality and the semantics of mood. In Searle, J., Kiefer, F., and Bierwisch, M., editors, *Speech Act Theory and Pragmatics*, pages 71–95. Springer.
- Kaplan, D. (1997). The meaning of 'ouch' and 'oops'. Ms., UCLA.
- Karlsun, F. (1999). *Finnish: an essential grammar*. Routledge.
- Karttunen, L. and Kay, M. (1985). Parsing in a word order free language. In Dowty, D., Karttunen, L., and Zwicky, A., editors, *Natural language parsing: psychological,*

- computational and theoretical perspectives*. Cambridge University Press.
- Kratzer, A. (1999). Beyond ‘ouch’ and ‘oops’: How descriptive and expressive meanings interact. Paper presented at the Cornell conference on theories of context dependency.
- Kratzer, A. and Shimoyama, J. (2002). Indeterminate pronouns: the view from Japanese. In Yukio, O., editor, *Third Tokyo Conference on Psycholinguistics*, pages 1–25. Hituzi Syobo, Tokyo.
- Krifka, M. (2001). Quantifying into question acts. *Natural Language Semantics*, 9:1–40.
- Krifka, M. (2014). Embedding illocutionary acts. In Roeper, T. and Speas, M., editors, *Recursion, complexity in cognition*, pages 125–155. Springer.
- Ladd, D. R. (1981). A first look at the semantics and pragmatics of negative questions and tag questions. *Chicago Linguistics Society*, 17:164–171.
- Lascarides, A. and Asher, N. (2009). Agreement, disputes and commitments in dialogue. *Journal of Semantics*, 26(2):109–158.
- Malamud, S. and Stephenson, T. (2014). Three ways to avoid commitments: declarative force modifiers in the conversational scoreboard. *Journal of Semantics*, pages 1–37.
- Matthewson, L. (2011). On apparently non-modal evidentials. pages 333–357.
- Merin, A. and Nikolaeva, I. (2008). Exclamative as a universal speech act category: a case study in decision-theoretic semantics and typological implications. Ms., University of Konstanz and SOAS London University.
- Michaelis, L. (2001). Exclamative constructions. In Haspelmath, M., editor, *Language typology and language universals: An international handbook*, volume 2, pages 1038–1050. De Gruyter.
- Murray, S. (2010). *Evidentiality and the structure of speech acts*. PhD Thesis, Rutgers University.
- Murray, S. (2011). A Hamblin semantics for evidentials. In Ito, S., Cormany, E., and Lutz, D., editors, *Proceedings of SALT 19 (2009)*. CLC.
- Murray, S. (2014). Varieties of update. *Semantics & Pragmatics*, 7:1–53.
- Murray, S. and Starr, W. (2016). The structure of communicative acts. Ms., Cornell University.
- Portner, P. (2004). The semantics of imperatives within a theory of clause types. In Young, R., editor, *Proceedings from Semantics and Linguistic Theory (SALT) 14*, pages 235–252. Cornell University Press.
- Portner, P. (2006). Comments on faller’s paper. Workshop on Philosophy and Linguistics at the University of Michigan.
- Portner, P. (2016). Commitment to priorities. In Fogel, D., Harris, D., and Moss, M., editors, *New work on speech acts*. Oxford University Press.
- Poschmann, C. (2008). All declarative questions are attributive? *Belgian Journal of Linguistics*, 22:247–269.
- Potts, C. (2003a). Expressive content as conventional implicature. In Kadowaki, M. and Kawahara, S., editors, *Proceedings of the North East Linguistic Society 33*, pages 303–322, Amherst, MA. GLSA.
- Potts, C. (2003b). A layered semantics for utterance modifiers. Ms., Brown Workshop on Direct Compositionality.
- Potts, C. (2005). *The logic of conventional implicature*. Oxford University Press.
- Potts, C. (2007). The expressive dimension. *Theoretical linguistics*, 33:165–197.
- Reese, B. (2008). *Bias in questions*. PhD Thesis, University of Texas, Austin.
- Rett, J. and Murray, S. (2013). A semantic account of mirative evidentials. In *SALT*

- XXIII Proceedings*, pages 453–472. CLC Publications.
- Roberts, C. (1996). Informative structure in discourse: towards an integrated formal theory of pragmatics. In Yoon, J. and Kathol, A., editors, *OSU Working Papers in Linguistics*, volume 49. OSU.
- Roberts, C. (2015). The character of epistemic modality: evidentiality, indexicality, and what’s at issue. Ms., Ohio State University.
- Romero, M. and Han, C.-h. (2004). On negative yes/no questions. *Linguistics and Philosophy*, 27:609–658.
- Rooth, M. (1985). *Association with Focus*. PhD Thesis, University of Massachusetts, Amherst.
- Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics*, 1:75–116.
- Rooth, M. (1996). Focus. In Lappin, S., editor, *The Handbook of Contemporary Semantic Theory*, pages 271–298. Blackwell.
- Sadock, J. M. (1974). *Towards a linguistic theory of speech acts*. Academic Press.
- Searle, J. and Vanderveken, D. (1985). *Foundations of illocutionary logic*. Cambridge University Press.
- Searle, J. R. (1969). *Speech acts*. Cambridge University Press.
- Simons, M., Tonhauser, J., Beaver, D., and Roberts, C. (2010). What projects and why. In Li, N. and Lutz, D., editors, *Proceedings of SALT 20*. CLC.
- Slobin, D. and Aksu, A. (1982). Tense, aspect and modality in the use of the turkish evidential. In Hopper, P., editor, *Tense-Aspect: Between semantics and pragmatics*, pages 185–200. John Benjamins.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2:447–457.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9:315–332.
- Stenius, E. (1967). Mood and language-game. *Synthese*, 17:254–274.
- Sudo, Y. (2013). Biased polar questions in english and japanese. In Gutzmann, D. and Gaertner, H.-M., editors, *Beyond Expressives: Explorations in Use-Conditional Meaning*. Brill, Leiden.
- Vallduví, E. and Vilkuna, M. (1998). On rheme and kontrast. In Culicover, P. and McNally, L., editors, *Syntax and semantics 29: The limits of syntax*, pages 79–108. Academic Press.
- Vanderveken, D. (1990). *Meaning and speech acts*. Cambridge University Press.
- Šafářová, M. (2007). Nuclear rises in update semantics. In Aloni, M., Buter, A., and Dekker, P., editors, *Questions in dynamic semantics*. Elsevier.
- Walker, M. (1996). Inferring acceptance and rejection in dialogue by default rules of inference. *Language and speech*, 39:265–304.
- Waltereit, R. (2001). Modal particles and their functional equivalents: a speech-act theoretic approach. *Journal of Pragmatics*, 33:1391–1417.
- Winans, L. (*to appear*). Disjunction and alternatives in Egyptian Arabic. In Onea, E., Zimmermann, M., and von Heusinger, K., editors, *Questions in Discourse*. Brill.
- Wu, J.-S. (2008). Antonyms? Presuppositions? On the semantics of two evaluative modals *jingran* and *guoran* in Mandarin. *Journal of Chinese Language and Computing*, 18:161–173.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116:983–1026.
- Zimmermann, M. (2011). Discourse particles. In von Heusinger, C., Maienborn, C., and Portner, P., editors, *Semantics*, pages 2011–2038. de Gruyter.