Intonation and Expectation: English Mirative Contours and Particles[1]

Kelsey KRAUS — *University of California Santa Cruz*

Abstract. This paper proposes a novel account of the English discourse particles *oh* and *huh*. First, these particles are analyzed as being inherently mirative, betraying a speaker's (violated) expectations in a discourse. Second, these particles are systematically paired with idiomatic prosodic contours and contexts to observe how their pragmatic interpretation shifts. Mirativity in English, it turns out, is not limited to particular particles, but can also be a property of prosodic contours. Sag and Lieberman (1975)'s Surprise-Redundancy Contour, is one of them. I propose that understanding the contribution of discourse particle requires intricate pragmatic reasoning strategies which enriches basic meaning of a discourse particle with the pragmatic discourse effects contributed by prosodic variation.

Keywords: discourse particles, prosody, mirativity, expectation, commitment.

1. Introduction

Discourse particles are pragmatic elements that signal a change in a speaker's knowledge state, used to help conversational participants navigate a discourse (Schiffrin, 1987). But this description is incredibly vague. Elements that belong to this class in English can mark acknowledgment (1a), betray a speaker's confusion (1b), or signal a change in topic (1c):

(1) Jeff made pizza dough this morning.
   a. **Oh**. Maybe he'll share some with us.
   b. **Huh**. I thought he didn't know how to cook.
   c. **Well**, I'm making brioche.

I restrict the focus here to two particles: *oh* and *huh*. Here, I ask first what effect these elements have on a larger discourse structure, and second, how much their *performance* impacts interpretation, as in (2):

(2) A: We're all out of flour.
   a. Oh. There's some in the pantry. + *H* L-L%* *Neutral Final Fall*
   b. Oh?! There's some in the pantry! + *(H) L* H* L-L%* *Surprise-Redundancy Fall*

I assume that the interpretation of discourse particles and intonation is pragmatic, and that both track participants' expectations in a discourse context. Such particles are very semantically bleached, and as such, their interpretation is especially sensitive to their prosodic environment. A thorough look at discourse particles must necessarily tease apart the pragmatic contribution of the particle from the interpretation of its prosodic context. With intonational and contextual environments held constant, we can identify places where apparently identical elements pull apart. *Oh* and *huh* appear to serve the same discourse function in (3), but changing the response type, as in (4), shows they are not:

[1]Thanks to Pranav Anand, Donka Farkas, Grant McGuire, Paul Willis, Deniz Rudin, Steven Foley, Tom Roberts, Netta Ben-Meir, and audiences at QSA17. This work is supported by the NSF GRFP Grant No. DGE 1339067.

(3) **A:** The server is down.
   a. **B:** Oh. They must be updating it.
   b. **B:** Huh. They must be updating it.

(4) **A:** The server is down.
   a. **B:** Oh. No it isn't.
   b. **B:** #Huh. No it isn't.

Deconstructing the pragmatic contributions of contours and particles allows us to make predictions about how particles will behave with particular contours, which make testable claims about the additivity of particles' and contours' contributions as speech-act-modifying elements.

2. Defining the Space

2.1. Theoretical Underpinnings: the Table Model of Commitment

Farkas and Bruce (2010)'s Table model of discourse (based on the works of Gunlogson, 2001; Stalnaker, 1978) is a system developed to model commitments, salience, and common ground management. This model decomposes the act of assertion into the effects that it has on a discourse that do not become part of the common ground, and seeks to unify the way in which declaratives and interrogatives are formally represented. Farkas and Bruce (2010) expand the original Stalnakerian model composed of a *Context Set* and a *Common Ground* and propose the addition of a stack of at-issue propositions called the *Table*, sets of *Discourse Commitments* indexed to particular speakers, and a *Projected Set* of future common grounds, accessible from the current contextual state. Assertion is then be defined by the effects that it has on these three elements of the model. The essential components are summarized in (5):

(5) a. **Table:** a stack of issues modeling salience in the current context
   b. **Common Ground (*cg*):** the set of all propositions that all members of the conversation are publicly committed to during the course of the conversation
   c. **Discourse Commitments (DCs):** for every participant x, the set of propositions that x has publicly committed to, but which are not yet part of the CG
   d. **Projected Set (*ps*):** the set of propositions currently under consideration for addition into the *cg*
   e. **Context Set (*cs*):** the set of all worlds that are compatible with the propositions in the Common Ground

Participants in a conversation move forward through a shared commitment to growing the *cg*, and shrinking the *cs*. This is done when speakers raise issues, place them on the Table for consideration, and resolve them by clearing them from the Table. Assertion is modeled by raising an Issue $\{p\}$, which places that issue on the table, which also updates a speaker's *DC*s with p. Farkas and Bruce (2010) define this in terms of a function from an input context K_i, in which an author a asserts a sentence s that denotes a proposition p, to an output context K_o:

(6) $Assert(a, p, K_i) = K_o$ such that: *Based on Farkas and Bruce (2010), p. 92*
   a. $\mathrm{DC}_{a,o} = \mathrm{DC}_{a,i} \cup \{p\}$
   b. $\mathrm{Table}_o = \mathrm{Table}_i + \{p\}$
   c. $ps_o = ps + \{\mathrm{cg}_i + \mathrm{p}\}$

Schematically, this can be visualized in (7), where an initial conversational state (7a) is updated by A's assertion (7b), which places $\{p\}$ on the table and p in A's *DC*s (7c).

(7) **a. Context before utterance:**

DC$_A$	**Table**	**DC**$_B$

$cg : s_0, ps = \{s_0\}$

b. A: asserts p

c. Update context with p**:**

DC$_A$	**Table**	**DC**$_B$
p	$\{p\}$	

$cg : s_0 = s_1, ps = \{\{s_0 \cup p\}\}$

This models the effect of an assertion, which is a a proposal to add the content of an utterance to the *cg*, and a speaker's public commitment to that proposition.

2.2. Expectation

I base much of this work on interpreting a speaker's subjective epistemic bias, which I term 'expectation'. This is influenced by Farkas and Roelofsen (2017)'s work on a speaker's credence level in p. I assume expectation for a proposition is based off of a speaker's (non-categorical) belief calculation for that proposition; expectation is a measure of relative credence levels.[2]

A speaker's expectation for a proposition or event denoted by q is built off her credence level in q, which is conditioned off of various other propositions in the *cs*. If a speaker believes q, we say that her expectation in q is relatively high, calculated against the other propositions in her *cs*. An agent's expectation for some q is based on the joint probability of that proposition and other *p*s in the *cs*. Expectation is high for q when all *p*s in the *cs* have a high joint probability with q. In contrast, a low joint probability between q and any context-set p should be enough to lower the expectation in q. This suggests the following definition for speaker expectation:

(8) Expectation for q by a speaker α:
$\text{Exp}_\alpha(q) = \min_{p \in CS} P_\alpha(p, q)$

A speaker α *believes* or *expects* q if $\text{Exp}_\alpha(q) \approx 1$. Subjective epistemic bias can be represented as an inequality in expectations for a proposition's truth. $\text{Exp}_\alpha(\varphi) > \text{Exp}_\alpha(\neg\varphi)$ is a way of talking about a speaker's non-categorical belief in φ, and will prove useful for talking about surprise, which can also be packaged as an instance of a speaker's violated expectations.

2.3. Calculating expectation: Mirative strategies

Many languages have grammaticalized the notion of surprise through particular marking on the verb, on nominals, or as freestanding morphemes, which is known as mirativity. I argue that a miratives are used primarily to communicate a speaker's expectation toward some event. DeLancey (1998); Aikenvald (2012) show through detailed investigation of many languages and language families that mirative strategies overtly map a speaker's surprised reaction toward a proposition to the current state of the common ground. These marked discourse moves add pragmatic information about a speaker's epistemic state to modify literal semantic contributions. The default state of the discourse assumes that there is some baseline range of expec-

[2] I use Expectation as a technical term, which does not necessarily carry the same intuitions as the word *expect*.

tation that most utterances will fall into, and as such, we need not mark violated expectations with every discourse move. Miratives convey that some eventuality falls outside of that normal baseline range; English mirative strategies express these ideas by way of discourse particles and intonation. Consider the difference between (9a) and (9b):

(9) a. Jeff ate the whole tub of yogurt. b. **Wow!** Jeff ate the whole tub of yogurt!

Assuming neutral intonation in (9a), we cannot read much into any pragmatic-level contribution of the speaker. But (9b) conveys an extra pragmatic effect: the speaker is surprised by the content of their utterance. Mirativity marks *how a speaker integrates new knowledge with old knowledge* in a discourse, and how that may affect the expectations she has established.

2.4. Representing Expectation in a Table Model of Commitment

Assertions require speaker commitment to the truth of an utterance. Whereas propositional content that is added to the table is a proposal that must be agreed upon by other discourse participants, the fine-grained pragmatic contributions of an utterance's performance is left out.

Imagine that we have three actors who would like to meet for dessert. Jeff and Tom are waiting for Sophie, and have decided to get ice cream. They notify Sophie and head out. Tom receives a reply from Sophie, informing him that she's lactose intolerant. Prior to this message, both Jeff and Tom assumed this choice was fine for everyone. But when Tom utters (10), the Table model cannot capture the pragmatic calculations that Jeff must compute to make sense of Tom's utterance. When Tom utters (10), he indicates his surprise at receiving this message, while also putting this proposition on the table for Jeff to consider:

(10) **Tom:** Sophie can't go to Mission Hill Creamery! $= p$
Extra inference: *Tom is surprised that p*

The table model as it stands does not have a way to model expectations, which are integral to how actors fit their contributions into a larger context. Intonation works at the speech act level, fine-tuning the exact sincerity contribution an utterance makes. When a speaker uses a neutral falling tone on a declarative, they present themselves as believing the content of that utterance. Doing this means that they must also *expect* this utterance to be true. In (10), Tom's tone is overall falling, but the final fall has a higher pitch excursion than normal, signaling to Jeff that he is asserting this proposition, as well as expressing surprise about it. That surprise is speaker-oriented, and it can't be questioned by the listener; it is infelicitous for Jeff to respond targeting Tom's surprise:

(11) **Jeff:** #You're not surprised about that.

One level of a speaker's *DC*s comes from the act of uttering semantic content, which places a proposition on the table and adds it to their *DC* list. I propose to separate the act of assertion from pragmatic contribution of prosody, logged as illocutionary not-at-issue-content is into the *DC* list. This information is discourse-relevant as commentary that the speaker makes public, but non-negotiable by other players. Using a neutral utterance-final falling tune (H* L-L%) commits the speaker to having a high expectation toward the truth of their utterance:

(12) The neutral H*L-L% contour adds the following to the speaker's DC:
$\text{Exp}_{spkr}(p) \approx 1$

Adding expectational content from intonation to the *DC* list allows us to easily model the pragmatic effect of intonation and discourse particles.

3. Contours and Particles

This work seeks to explain the effects of three English mirative strategies: the SRC, *oh*, and *huh*. Though superficially different, these markers have the following four things in common:

(13) a. *Optionality*: never semantically necessary, only contribute pragmatic content
b. *Illocutionary not-at-issueness*: updates are to a speaker's expectational state
c. *Speech act modifiers*: comment on elements above the propositional level
d. *Contextual dependence*: use is interpretable only in context

I argue that these elements are inherently mirative, commenting on a speaker's (violated) expectations in context. The following sections propose that interpretations of the Surprise-Redundancy Contour, the particle *oh* and the particle *huh* can be analyzed along similar grounds.

3.1. Surprise-Redundancy Contour

Sag and Lieberman (1975) identify the surprise-redundancy contour (SRC) as an utterance-level prosodic tune that carries pragmatic information about speaker attitude. Their classification assumes that the contour is an intonational idiom consisting of a high pitch anchored to the primary sentential accent, and contrasting low pitch on the utterance's second most prominent syllable left of the main stress placement, as in (14) (Hayes, 1995; Ladd, 2008):

(14) a. [T h e blackb oar d's pa in te d o r a n g e !]
H L* H* L

The blackboard's painted orange!↘

b. [Alic e is n't c o m i n g !]
(H)L* H* L

Alice isn't coming!↘

Sag and Lieberman (1975) identify this contour as expressing a speaker's surprise at a proposition or event in the world, much as a speaker might do in (14a) upon walking into a classroom with a brightly painted chalkboard, or in (14b) when, exasperated, they must repeat something they believes the hearer "should have known". If it is established that Alice is currently vacationing in Bermuda, asking why she isn't around is redundant. Using the SRC in 'redundancy' cases is akin to overtly questioning your addressee's information state. For a detailed look at the differences between 'surprise' and 'redundancy' readings of the SRC, see Kraus (2018).

Though it requires two stressed tones, the SRC can also appear on prosodic phrases with a single prominent position, or even a single syllable. To do this, the vowel is lengthened, giving the effect of two syllables:

(15) a. Mother: Did you brush your teeth?
Child: Du- uh!
L* H*-L%

b. Mother: Don't forget to floss!
Child: Mo- om!
L* H*-L%

A speaker uses the SRC to point to something about the a participant's actions or utterances in previous discourse that they are not currently taking into account. With the SRC, a speaker highlights an issue that she believes is inferable, but is being ignored. I formalize this in (16):

(16) The English Surprise Redundancy Contour (H) L* H*-L% is anaphoric to a salient proposition or event p in a discourse context C, and is admissible for discourse-salient participants x when
   a. q is the proposition expressed by the speaker (uttered content or the presuppositions introduced by a question),
   b. add the following to the speaker's *DC*s:
   $\text{Exp}_{spkr}(q) \approx 1 \wedge \forall x \in C\ [\text{Exp}_x(p|q) \approx 0]$

I discuss the pragmatic contribution of the SRC more fully in Section §4.2.

3.2. English *oh*

Abstractly, *oh* marks some change of state that has occurred, whether in the speaker's reasoning or in their accessible information (Aijmer, 2002). (17) shows this particle used similarly to a response particle, whereas (18) is closer to a discourse particle use:

(17) *Speaker walks into a room full of snakes:*
Oh!

(18) *Speaker receives news that the room full of snakes is to be incinerated:*
Oh! That's great news!

Surprise can be captured in terms of violated expectations, marking a change in a speaker's reasoning from a previous state. Other emotions can also be thought of in terms of a change of state. Anger, for example, can be framed as a negative change in speaker expectations. When using *oh*, a speaker indicates both her own change of state, as well as primes the listener for an upcoming change to their own discourse model (Fox Tree and Schrock, 1999).

Gunlogson (2008) notes that *oh* can be used to diagnose speaker commitment when paired with a neutral fall(H* L-L%). A neutral *oh* response to a previous utterance has a similar effect to the response particle *yes*, committing the speaker to the content of the previous proposition. But unlike *yes*, *oh* signals that this is new information to the speaker:

(19) A: They quarantined that room.
   a. B: Yes. *commits to A, A can already be known to B*
   b. B: Oh. *commits to A, A is not already known to B*

In both cases, the particles commit the speaker to the content of A's utterance. Yet while *Yes* with rising intonation commits the speaker to the content of A's utterance, the same rise on *oh*

does not commit B to A's utterance (for an in-depth discussion of the varied contributions of types of rising intonation, see Jeong, 2017):

(20) A: They quarantined that room.
  a. B: Yes? *B is committed to content of A*
  b. B: Oh? *B is **not** committed to content of A*

Oh's contribution cannot be commitment. This is clear in out-of-the-blue cases, where a speaker reacts to some environmental factor, rather than a previous utterance:

(21) *Speaker walks past a charred building:*
  a. Oh! b. Oh! They burned the snakes!

Oh with excited intonation highlights a speaker's surprise about an event. While a follow-up utterance may provide some proposition to commit to, the particle in itself does not. (21a) simply conveys a speaker's violated expectations. In keeping with 'change-of-state' as a core meaning component of *oh*, I propose that the particle requires an inequality between a speaker's expectations and the information that they have been presented with, allowing the hearer to infer that something did not accord with how the speaker imagined the conversation would progress:

(22) *oh* is anaphoric to a proposition p salient in the discourse s.t.:
  a. *oh*(p) adds the following to the speaker's *DC*s:
  $\text{Exp}_{spkr}(p) < \text{Exp}_{spkr}(\neg p)$

Oh indicates that the speaker had higher expectations for $\neg p$ over p. Expectation allows the contribution of *oh* to operate on a pragmatic level, commenting on a speaker's view with immunity to public scrutiny as shown in (23):

(23) A: The University has banned reptile rooms.
  B: Oh.
  A: #That's not true. You expected that.

Formulating the contribution of *oh* in terms of expectations casts this particle as a mirative marker. Mirativity relies on the reinterpretation of a speaker's expectations. With excited intonation, the mirative strategies have an additive effect. In an utterance like (21), the speaker can augment her degree of surprise by using an excited contour. We will come back to this point, as the interaction of particles and intonation will be crucial to understanding their overall discourse function.

3.3. English *huh*

In many cases, *huh* patterns alongside *oh*, leading many to assume them to be interchangeable. Indeed, it is also a particle that indicates a shift in speaker expectations:

(24) *Speaker receives news that the room full of snakes is to be incinerated:*
  a. Huh! b. Huh! That's great news!

But in particular cases, *huh* patterns differently than *oh*. One place is in cases where two participants have equal information. In (25a), B can use *oh*, although she overtly indicates that she is already aware of the content of A's utterance, but she can not indicate this with *huh*.

(25) They quarantined that room.
   a. B: Oh. I know.
   b. B: #Huh. I know.

Huh's behavior also patterns differently in the Gunlogsonian test for commitment to a previous utterance. B can utter (26a) without going to check on the status of the room in question. But uttering (26b) is odd if B has not gotten up, and verified A's assertion:

(26) A: That room is full of snakes.
   a. Oh. No it isn't. (It just looks like that.)
   b. #Huh. No it isn't.

In (26b), the particle and the continuation utterance are two unconnected conversational moves, whereas (26a) serves as a reaction and an explanation for the reaction. It appears that *huh* must acknowledge and at least temporarily accept the previous proposition. I propose a formulation of *huh*'s discourse contribution in much the same way as *oh*. *Huh* also requires a contextual inequality between a speaker's expectations and the information that she is faced with, as well as a further restriction.

(27) *huh*(p) adds the following to the speaker's *DC*s:
$\text{Exp}_{spkr}(p) < \text{Exp}_{spkr}(\neg p) \wedge \text{Exp}_{spkr}(p) > 0$

Huh requires that the speaker have expected $\neg p$ over p while also commenting on the speaker's relative expectations about p. This is clear with contextual differences in cases of speaker surprise. With an excited contour, out-of-the-blue scenarios of surprise are bad with *huh*:

(28) *Speaker, oblivious, rounds the corner and is hit with confetti:*
   a. Oh! I didn't expect that!
   b. #Huh! I didn't expect that!

But if a surprise is expected, for example, in the case that the speaker and her colleagues occasionally play practical jokes, both *oh* and *huh* are available:

(29) *Speaker, an occasional practical joker, rounds the corner and is hit with confetti:*
   a. Oh!/Huh! I didn't expect that!

Again, formulating *huh*'s contribution in terms of violated expectations has the advantage of comparing this particle to other markers of mirativity.

4. Compositionality: the relation of text to tune

Previous accounts of English discourse particles show that these elements are impressively flexible, but no research to date looks at the compositionality of discourse particle meaning and intonational meaning. In an effort to sharpen the picture even further, I outline a number of contexts that turn out to have an interesting effect on a particle performance's felicity. By

pinpointing specific discourse contexts and holding them constant, we can see what other kinds of extra-linguistic factors these pragmatic markers are sensitive to. The felicity of these particles is not limited to their combinability with particular contours; they interact with the overall structure of the discourse, too. Due to space restrictions, I present a subset of relevant discourse contexts here to show this difference. For a fuller picture, see Chapters 3 & 4 of Kraus (2018), which lays out a more in-depth theoretical argument backed up by experimental evidence.

Below is an overview of four of the relevant contexts and manipulations, and the behavior of *oh* and *huh* with each context + prosody pair:

	Neutral Fall H* L-L%		**SRC Fall (H) L* H* L-L%**	
Context	*oh*	*huh*	*oh*	*huh*
Eavesdropping	✓	✓	✓	✓
Sudden Realization	✓	✓	*	✓
Contradict a Statement	✓	*	*	✓
Correct a fact	✓	*	✓	✓
Accept a fact	*	*	✓	✓

Table 1: Contours, contexts, and particle combinations

In broad strokes, when paired with the neutral falling contour, *huh* is licensed in a proper subset of the environments that license *oh*. This generalization is reversed when *oh* and *huh* appear with the SRC. In all these cases, the additive properties of a particle's pragmatic contribution and a contour's contribution can derive the patterns seen in the table above. The following two sections break this down and show just how we can use speaker expectation to understand how these patterns emerge in each discourse context.

4.1. Text and Tunes: Neutral Falling H*L%, *oh*, and *huh*

Recall that *oh* and *huh* differ only in the expectations that are placed a speaker's *DC* list:

(30) *oh* and *huh* are anaphoric to a proposition p salient in the discourse s.t.:
   a. *oh*(p) adds the following to the speaker's *DC*s:
   $\text{Exp}_{spkr}(p) < \text{Exp}_{spkr}(\neg p)$
   b. *huh*(p) adds the following to the speaker's *DC*s:
   $\text{Exp}_{spkr}(p) < \text{Exp}_{spkr}(\neg p) \wedge \text{Exp}_{spkr}(p) > 0$

The difference between the two is the extra pragmatic effect that *huh* expresses. In contrast to *oh*, when using *huh*, the speaker conveys that she did not entirely rule out the situation that she currently finds herself in. In what I assume is the base case, neutral, falling contours on discourse particles have the effect of an assertion—a property of an intonational contour, not of a particular sentence form. I assume that neutral, falling contours assert the content of the proposition they take as their prejacent. I follow Farkas and Roelofsen (2017) and Malamud and Stephenson (2015) in assuming that a falling declarative asserts speaker belief in p, which I formulate in terms of expectation contributed by the contour:

(31) The neutral falling H*L-L% adds the following to the speaker's *DC*s:
$\text{Exp}_{spkr}(\text{p}) \approx 1$

For the most part, *oh* and *huh* pattern in very similar ways, and are available in overlapping contexts. Both can be used with the neutral contour to indicate sudden realization (32) and eavesdropping situations (33), and are both disallowed in cases of fact acceptance, as in (34):

(32) *Speaker walks out of a building, with no expectation of the weather.*
Oh./Huh. It's raining. *Sudden Realization*

(33) *Characters are speaking to other characters in a film.* A: It's Tess Ocean!
B, *also observing the film:* Oh./Huh. It's clearly Julia Roberts. *Eavesdropping*

(34) A, (*Quizmaster*): What's the capital of Delaware?
B: Dover.
A: It is Dover.
B: #Oh./#Huh. I was right. *Accept a fact*

Oh and *huh* do pull apart when contextual assumptions shift. The Gunlogsonian test for commitment with the neutral falling contour works with *huh* in cases where B has less evidence or a lower expectation than A (35). But when a speaker has equal to or higher expectation for the truth of a proposition, *huh* is no longer a licit response, and is also no longer a diagnostic for commitment (36):

(35) A: The server's down.
a. B: Oh. (I didn't know that.)
b. B: Huh. (I didn't know that.)

(36) A: The server's down.
a. B: Oh. I know.
b. B: # Huh. I know.

If B uses *huh* in (36), they set up an inconsistent set of *DC*s for them to go on to inform A that this issue is already known. *Huh* reacts to the speech event here: B is surprised to be informed of something is believed to be common ground. Using *huh* communicates she is communicating the following things about her doxastic state:

(37) B: # Huh. I know.
a. B did not expect A to assert p, "The server is down" $[\text{Exp}_A(p) < \text{Exp}_A(\neg p)]$
b. B holds a non-zero expectation for the possibility that A will make the speech event "The server is down."

While this is a consistent belief state, it is contradictory to convey this to an addressee. When surprise is directed at the level of illocution, *huh* is a strange pragmatic choice. This seems to be a special case of fact contradiction contexts, as in (36), repeated in (38) where a speaker attempts to use *huh* in an instance where they mean to directly contradict the previous statement:

(38) A: That room is full of snakes.
a. B: Oh. No it isn't. (It just looks like that).
b. B: # Huh. No it isn't.

This fact is borne out in the formulations of these particles where *huh* has an added restriction that there must be some expectation already calculated to a given surprising proposition.

This predicts that *oh* is more permissive than *huh*, particularly in cases of extreme surprise or with events where the speaker may have had zero expectation for p. For *huh*, while the context does favor expectation for $\neg p$ over p, it still requires p to be an available option for the speaker's chain of reasoning.

Both particles are mirative marking strategies, which rely on the expectations of the speaker. In most of the contexts discussed so far, the surprising element is an issue raised and placed on the table by an interlocutor. But as we have just seen, out of the blue cases of surprise, as well as some correction cases, target a larger event. This ability to target illocutionary content is part of the mirative function of these particles, coupled with pragmatically introducing a speaker's own subjective bias. Placing a speaker's expectational commitments in a table model of discourse is a convenient way of keeping track of issues raised, common ground material, and conversational participants' subjective epistemic states. For particle only responses like (35), a speaker's discourse commitments are updated as in (39):

(39) A: The server is down.

a. Context after A's utterance:

DC_A	Table	DC_B
p, $\text{Exp}_A(p) \approx 1$	$\{p\}$	

$cg : s_0$, $ps = \{\{s_0 \cup p\}\}$

b. B utters *huh*

B: Huh.

c. Context after B utters *huh*:

DC_A	Table	DC_B
p, $\text{Exp}_A(p) \approx 1$	$\{p\}$	$\text{Exp}_B(p) < \text{Exp}_B(\neg p)$ $\wedge\ \text{Exp}_B(p) > 0$

$cg : s_0$, $ps = \{\{s_0 \cup p\}\}$

With no clear commitment one way or another, the speaker and hearer may choose to take this particle response as committing B to p, thus clearing the scoreboard. This reasoning is due in part to the contribution of the final fall. In (39), a neutral falling contour on the particle allows A to infer that B has accepted this contribution. This can then grow the common ground. When B responds with the particle and a following utterance, the expectations that follow are logged in their *DC*s, and the propositional content of p is placed on the Table. Pragmatic reasoning takes the entire utterance and performance into account to determine whether B's contribution was an accepting or rejecting move.

I modify the Table model to include the stack of propositions on the table that have yet to be accepted or rejected. In (40), neither p nor q has been explicitly accepted by both participants. For the sake of simplicity, I assume that only the most recent proposition in the stack, q is represented in the projected set, though there may be other approaches to representing this state of a discourse.

(40) A: The bank is closed today. = p
B: Oh. I need to deposit a check. = q

a. Context after A's utterance:

$\mathbf{DC}_A$	**Table**	$\mathbf{DC}_B$
p, $\text{Exp}_A(p) \approx 1$	$\{p\}$	

$cg : s_0$, $ps = \{\{s_0 \cup p\}\}$

b. B utters *oh* + q

c. Context after B utters *oh*+q:

$\mathbf{DC}_A$	**Table**	$\mathbf{DC}_B$
p, $\text{Exp}_A(p) \approx 1$	$\{\text{q}\}$ $\{\text{p}\}$	q, $\text{Exp}_B(q) \approx 1$, $\text{Exp}_B(p) < \text{Exp}_B(\neg p)$

$cg : s_0 = s_1$, $ps = \{\{s_0 \cup q\}\}$

B's utterance here is not a clearly accepting or rejecting move. As *oh* does not supply this confirmation, and B's following utterance leaves the option open, both p and q are issues up for discussion. We reason that the falling contour on both *oh* and q lead A to infer that p is resolved, and cleared from the table. I assume that the individual discourse commitments of a particular speaker are related by conjunction. Pragmatic meanings contributed by particles and contours are calculated additively by the expressions in a participant's discourse commitments.

This distinction between non-zero and no expectation that drives the differences between *oh* and *huh* can explain the oddness introduced by fact correction scenarios, as in (41):

(41) A: Sandy is from Nebraska.

a. B: Oh. She's from California.
H* L-L%

b. B: # Huh. She's from California.
H* L-L%

In both (38) and (41), B disagrees with A's proffered content, in the former, by offering explicit contradiction, and the latter, by offering up conflicting information that B knows. In both of these scenarios, *oh* is allowable, while *huh* is not. Knowing a fact q about the world implies high expectation for it to be true: $\text{Exp}_{spkr}(q) \approx 1$. If someone utters p, a fact that would make q false, a speaker can signal this discrepancy with *oh*; based on their belief in q, they have no reason to believe p to be true. If B knows that Sandy is from California, he has reason to believe that she is not from other places. With this information in B's discourse commitments, but not actively on the Table, the listener can infer the following:

(42) *Sandy is from Nebraska* = p *Sandy is from California* = q

a. Context after A's utterance:

$\mathbf{DC}_A$	**Table**	$\mathbf{DC}_B$
p, $\text{Exp}_A(p) \approx 1$	$\{p\}$	

$cg : s_0$, $ps = \{\{s_0 \cup p\}\}$

b. B utters *huh* + q

c. Context after B utters *huh* + q:

$\mathbf{DC}_A$	**Table**	$\mathbf{DC}_B$
p, $\text{Exp}_A(p) \approx 1$	$\{\text{q}\}$ $\{\text{p}\}$	q, $\text{Exp}_B(q) \approx 1$, $[\text{Exp}_B(p) <$ $\text{Exp}_B(\neg p) \wedge \text{Exp}_B(p) > 0]$

$cg : s_0 = s_1$, $ps = \{\{s_0 \cup q\}\}$

Pragmatic calculation:

i. **Falling contour:** $\text{Exp}_B(q) \approx 1$

ii. ***huh*:** $\text{Exp}_B(p) < \text{Exp}_B(\neg p) \wedge \text{Exp}_B(p) > 0$

iii. **Pragmatic inconsistency:** $q \rightarrow \neg p$, so $\neg p = 1$. Since q and p are contrary, one cannot commit to q being true and to p being possible.

Huh's added restriction that the expectation for p be non-zero leads to an inconsistent pragmatic calculation. If B knows his information is correct, there is no way that he can felicitously signal that with *huh*. But he can with *oh*:

(43) *Sandy is from Nebraska = p* *Sandy is from California = q*
   a. **Falling contour:** $\text{Exp}_B(\text{q}) \approx 1$
   b. ***oh*:** $\text{Exp}_B(p) < \text{Exp}_B(\neg p)$
   c. **Pragmatic consistency:** $q \rightarrow \neg p$, so $\neg p = 1$. B must have expected p to be 0.

Notice how *huh* again becomes licit when B's uncertainty is put back into play:

(44) A: Sandy is from Nebraska.
   a. B: Oh./Huh. I thought she was from California.

When B indicates overtly that they are not sure about the contribution, *huh* again becomes licit. The semantic content of this utterance allows for felicitous pragmatic interpretation of *huh*.

4.2. Text and Tunes: Surprise-Redundancy, *oh*, *huh* and *what*

The previous sections showed the overlap in meaning between *oh* and *huh* in the pragmatic context of a neutral final fall. In this section, the pragmatic effects of the SRC show a different split between the conditions of use of *huh* and *oh*.

Recall that SRC is more complex than other final falls, requiring at least a rise-fall pattern with two different prominent tone heights, (H) L* H* L-L%. The pragmatic effect of the SRC is complex as well–it indicates that the speaker believes the listener "should have known" some salient proposition in the discourse. This contour bakes in a mirative meaning itself, anchoring itself to speaker expectations about some participant in the discourse. Formally, I represent the pragmatic contribution of the SRC to the discourse in (45), repeated from (16):

(45) The SRC (H) L* H*-L% is anaphoric to a salient proposition or event p in a discourse context C, and is admissible for discourse-salient participants x when:
   a. q is the proposition expressed by the speaker (uttered content or the presuppositions introduced by a question),
   b. add the following to the speaker's *DC*s:
   $\text{Exp}_{spkr}(q) \approx 1 \wedge \forall x \in C\ [\text{Exp}_x(p|q) \approx 0]$

This formalization does not presume to break down the prosodic components of the SRC into meaningful sub-phrasal tones. Rather, it assigns a conventional pragmatic effect to this prosodic phrase-level melody. The SRC anchors speaker expectation to a proposition q, while expressing the speaker's secondary expectation that given q, everyone in the discourse context should have very low, if not zero expectation for p. In other words, the speaker expresses surprise that the conversational participants have deduced p, given what she expects them to have predicted by having access to q. The speaker does not expect p, and assumes no one else should either.

A small note on the notation used here to show the relationship between p and q is in order. The formulation $\text{Exp}_x(p|q) \approx 0$ could either refer to conditional probability or likelihood ratios. This should be understood in the likelihood function use as we are not making predictions about

some future outcome or state of the discourse, which expectation calculations based on conditional probabilities would assume. Because this contour is used in response to an unexpected or surprising event, the speaker is using this contour in the face of specific, observed data. We want to express that given the situational facts and observations, the computed likelihood for some at-issue proposition is quite low.[3]

Just like other miratives, using the SRC is a strategy that publicizes the speaker's current epistemic state. By placing its contributions in a speaker's *DC*s, a speaker signals to the listener that this is not at-issue content, but that it is content that the speaker is committed to:

(46) A, *(with SRC:)* This cheddar cheese is orange!↘
B: # That's not true. You aren't surprised at all.

The SRC can combine with both discourse particles of interest. So as not to confuse this contour with a rising declarative, in the examples below, I use a question mark followed by a fall '?↘' to represent an intended instance of the SRC. The following contexts show where *oh* and *huh* pattern the same with the SRC:

(47) *Character speaking to other characters in a film* A: It's Tess Ocean! *Eavesdropping*
a. B, *Moviegoer, observing the film:*
Oh?↘/Huh?↘ It's clearly Julia Roberts?↘

(48) A, *Quizmaster:* What's the capital of Delaware? *Fact Acceptance*
B: Dover.
A: It is Dover.
a. B: Oh?↘/Huh?↘ I was right!

Both particles interacts with the SRC in a subtly different way. With *oh*, there is a pragmatic inference that the addressee's utterance is very speaker-new. The speaker is taken off guard; they have not previously entertained the proposition in question. This nicely falls out from *oh*'s characterization in (30a), as no imposed floor on a speaker's expectation for a preceding utterance allows for a genuine violation of expectations.

Like *oh*, *huh* also express heightened levels of disbelief when paired with the SRC. Another contribution of this contour comes at the perlocutionary level: the contour not only expresses

[3]I choose to use likelihood ratios for a few reasons. One salient reason is that when using simple probability, average propositions and their negations will not have roughly equal probability. This is based off the observation that for a proposition p to hold, all of its entailments must also hold. But for $\neg p$ to hold, just one of p's entailments can be false. Given this, it will tend to be the case that out of p and $\neg p$, one of these will have an intrinsically higher base probability. But though this is valid, this fact should not have an impact on the calculation of surprise.

If we think about surprise in terms of joint, conditional, or even just raw probability, we miss the relevant idea of surprise, as we cannot relativize this to a particular context. We want to get rid of the idea that surprise can be calculated context-free, which is what we would predict if we used any of the measures above. Using a likelihood ratio of the joint probability of the antecedent and the at-issue proposition, normalized against the probability of the antecedent and the at-issue proposition ($\frac{P(p,q)}{P(p)\times P(q)}$), we no longer have a sensitivity to the number of entailments that are carried either by the antecedent (p) or the at-issue proposition (q). Likelihood ratios are a way of achieving the normalization that is needed for context sensitivity by unlinking violated expectations from raw probability.

a speaker's expectations, but also persuades the addressee to react. This contour is inherently a request for the addressee to bolster their claims before the speaker will commit to their discourse move. This effect is strong—so strong that it is strange for a discourse to end with an utterance that carries the SRC. The pressure is high for the addressee to respond, likely due to *huh* explicitly encoding that an alternate proposition is more likely.

The prosodic environment imposed by the SRC flips some of the patterns of grammaticality between *oh* and *huh*. Whereas composing *huh* with either the neutral or the excited contours revealed that this particle occurred in a subset of the contexts where *oh* could occur, the findings are switched with the SRC. *Oh* becomes the particle with the more limited distribution, which is a subset of the contexts in which *huh* can appear. One place where the particles behave differently is in fact correction and contradiction contexts. Recall that *huh* is disallowed in both of these contexts when paired with a neutral H* L-L% contour:

(49) A: Sandy is from Nebraska. *Fact Correction*
   a. B: Oh?↘ She's from California?↘
   b. B: Huh?↘ She's from California?↘

(50) A: That room is full of snakes. *Contradiction*
   a. B: #/? Oh?↘ No it isn't?↘
   b. B: Huh?↘ No it isn't?↘

One first thing to observe is that none of the B responses completely settle the issue at hand. Rather, they invite A to weigh in on the matter again, whether it be to provide counter-evidence, or to commit to the proffered alternative. This is a pragmatic effect from the SRC, calculated when the speaker projects their own relatively high expectations about the prejacent and the expectations generated about other participants' beliefs about the alternative. The SRC must comment on something mutually manifest. In the case that the speaker offers an alternative to *p* that their conversational participant *couldn't have been aware of*, the result is an incoherent discourse state, regardless of the particle used:

(51) *Sophie and Tom are supposed to pick Jeff up from the airport at 2pm. Tom, but not Sophie, receives a flight alert, telling him that Jeff will be delayed until 4pm. Later, they are talking about when the need to leave to pick Jeff up*:
Sophie: I can be ready to leave for the airport at 1:30pm.
Tom: # Oh?↘/Huh?↘ His flight doesn't come in until 4?↘

Though the intuition is not as sharp as in other cases, the felicity of *oh* is questionable in cases of direct contradiction. Recall though, the difference between *oh* and *huh* in neutral contexts. There, *huh* was illicit, but in an interesting way. It necessitated extra accommodation on the part of the speaker in order for it to be interpretable; *huh* responses in those cases constituted two conversational moves: a discourse particle response, an intermediate verification step, and then a second corrective conversational update. A similar thing happens with SRC *oh* responses to contradiction environments. (50a) is only licit if *oh* is one discourse move, followed by a verification, and a corrective follow-up utterance.

Oh is infelicitous here due to the SRC's discourse management strategy that allows a speaker to assert an alternative proposition while implicitly requesting the addressee's feedback on it.

The alternative that the speaker offers must be an alternative that the listener could arrive at by carefully considering the *cg* and the *ps*. *Huh* is no problem because its meaning bakes in a non-zero chance that an alternative explanation is available. Cases of contradiction or fact correction are not purely contradiction because the mirative component of *huh* assumes that however unlikely p may be in the speaker's estimation, it is still a viable alternative.

But *oh* has no requirement that the speaker's expectation of the prejacent be greater than 0. *Oh*'s freedom is a weak restriction: because there are alternative particles that require uncertainty about a proposition, using a particle that does not have this requirement confronts the listener with a potential Maxim of Quantity violation. The speaker uses this because they want to indicate that they expect p to be 0 (from 50):

(52) *This room is full of snakes* = p *The is room is not full of snakes* = q

a. Context after A's utterance:

$\mathbf{DC}_A$	**Table**	$\mathbf{DC}_B$
p, $\text{Exp}_A(p) \approx 1$	$\{p\}$	

$cg : s_0$, $ps = \{\{s_0 \cup p\}\}$

b. B utters *oh* + q

c. Context after B utters *oh* + q:

$\mathbf{DC}_A$	**Table**	$\mathbf{DC}_B$
p, $\text{Exp}_A(p) \approx 1$	$\frac{\{q\}}{\{p\}}$	q, $[\text{Exp}_B(p) < \text{Exp}_B(\neg p) \wedge \text{Exp}_B(p) > 0]$, $[\text{Exp}_{A,B}(p\|q) \approx 0]$

$cg : s_0 = s_1$, $ps = \{\{s_0 \cup q\}\}$

Pragmatic calculation:

i. **SRC:** $\text{Exp}_B(q) \approx 1 \wedge \forall x \in DC\ [\text{Exp}_{A,B}(p|q) \approx 0]$
ii. ***oh*:** $\text{Exp}_B(p) < \text{Exp}_B(\neg p)$
iii. **Pragmatic consistency (?):** $q \rightarrow \neg p$. Speaker believes q = 1, infer $\neg p$ = 1. There is no restriction that p be 0.
BUT: Pragmatic competition w/ *huh* can leads one to wonder why *oh* was used.

If the speaker wants to directly contradict p, then they must have evidence against p. But the SRC still allows for the possibility that the speaker's expectations might be violated, which sets up a (defeasible) inconsistency: a speaker can fully expect p to be false (from *oh*) but when they do this, they cannot also expect it to be *almost always* be false (from the SRC). This appears to be a by-speaker pragmatic calculation in need of empirical validation. While some argue that *oh* with the SRC in contradiction cases is pragmatically fine, the same individuals agree that there is a cline of acceptability in the following responses to the same utterance, with direct contradiction cases more pragmatically odd:

(53) A: Sandy is from Nebraska.
a. B: $^{??}$Oh?↘ No she isn't! She's from California!↘
b. B: $^{?}$Oh?!↘ She's from California!↘
c. B: Oh?!↘ She has a California driver's license?!↘

With *huh*, the listener accommodates that they are not being directly contradicted: the speaker's expectation for p was greater than zero, leaving room that expectation of q, however strong, is not certain. Though the SRC signals to the listener that by the speaker's estimation, they should have known an alternative proposition q, it invites further commentary from the listener

by leaving the door open for follow-up. A more robust pattern of acceptability emerges from cases of sudden realization:

(54) *Speaker walks out from a windowless building, having no expectation for it to be raining:*
a. #Oh?!↘ It's raining?!↘ b. Huh?!↘ It' raining?!↘

If a speaker is talking to herself, (54a) is odd. But when sudden realization cases involve another participant, *oh* gets better:

(55) *A and B are solving math problems. A turns to write one on the board and picks up the chalk in her left hand*:
B: Oh?!↘ You're left-handed?!↘

If a speaker is addressing herself, a Quantity implicature arises from the interaction of the SRC, the requirements of *oh*, and the fact that the speaker did not use a rising declarative, indicating that they can not commit to p. The listener assumes that the speaker must have had some reason for using *oh* when *huh* was available as well. They deduce that the speaker expects q to be true:

(56) *The fact that it's raining* = p *It is raining* = q
Context after B utters *oh* + q:

DC_A	Table
q, $[\text{Exp}_A(q) \approx 1$	{q}
$\wedge\ \text{Exp}_{A,B}(p\|q) \approx 0]$	{p}

$cg = \emptyset$, $ps = \{\{p\}, \{q\}\}$

Pragmatic Calculation:
i. **SRC:** $\text{Exp}_B(q) \approx 1 \wedge \text{Exp}_B(p|q) \approx 0$
ii. ***oh*:** $\text{Exp}_B(p) < \text{Exp}_B(\neg p)$
iii. **Pragmatic inconsistency:** $q = \neg p$. Speaker believes q = 1, infer $\neg p$ = 1.
There is no restriction that p be 0.
The speaker is both the addressee and the source of q; they hold one belief about the actual facts in the world, p, and another about their expectations, $q = \neg p$.

When p turns out to be true, the free variable in the definition of the SRC allows for the speaker to believe $q \approx$1, while also strongly expecting that p is not the case. But the speaker has also just observed that it is raining. Using a rising declarative has the pragmatic discourse effect of indicating the speaker cannot commit to p, the highlighted alternative. But facts about the world prove otherwise. The speaker has seen the rain, and knows p is true. Her internal state is in conflict: she cannot know for sure that it's raining (p=1), expect that it's not ($q = \neg p$=1), and expect that given q, the expectation for p should be 0. In this case, there appears to be such a thing as too much mirativity.

5. Conclusion

This paper serves as both an argument recognizing that English has active mirative strategies, as well as a petition for an updated look at how the contribution of discourse particles and

prosodic meaning can shape a conversation. Discourse particles contribute information about how a speaker has structured her expectations, and how those expectations have come to be met or violated. Intonation plays an extremely similar role. While some utterance-level contours contribute straightforward pragmatic effects of authoritativeness or excitement on the part of the speaker, others are more nuanced. Viewing these contours in terms of a speaker's expectations about a discourse has the advantage of pinpointing which contours convey authoritativeness by virtue of assigning the truth of a proposition a high expected value. It can also capture the effects of contours that express violated expectations in view of a discourse context. By reshaping the way we view intonation and discourse marker meaning, we create a common foundation for analyzing these two distinct systems. Doing this gives us a coherent way forward in interpreting the many ways a speaker sheds light on her internal belief structure. This in turn provides a glimpse of the conversational structure, and a way to talk about the interactions between the text of a sentence and its tune.

References

Aijmer, K. (2002). *English Discourse Particles: Evidence from a corpus*. John Benjamins.

Aikenvald, A. (2012). The essence of mirativity. *Linguistic Typology 16*, 435–485.

DeLancey, S. (1998). Mirativity: the grammatical marking of unexpected information. *Linguistic Typology 1*, 33–52.

Farkas, D. and K. Bruce (2010). On reacting to assertions and polar questions. *Journal of Semantics 27*, 81–118.

Farkas, D. and F. Roelofsen (2017). Division of labor in the interpretation of declaratives and interrogatives. *Journal of Semantics*, 1–53.

Fox Tree, J. E. and J. C. Schrock (1999). Discourse markers in spontaneous speech: Oh what a difference an oh makes. *Journal of Memory and Language* (40), 280–295.

Gunlogson, C. (2001). *True to form: rising and falling declaratives as questions in English*. Ph. D. thesis, UC Santa Cruz.

Gunlogson, C. (2008). A question of commitment. *Commitment: Belgian Journal of Linguistics 22*, 101–136.

Hayes, B. (1995). *Metrical Stress Theory: Principles and case studies*. Chicago: The University of Chicago Press.

Jeong, S. (2017). Intonation and Sentence Type Conventions: Two types of rising declaratives. *Presented at the UCSC/Stanford Workshop on Clause Types, 29 January 2017*.

Kraus, K. (2018). *Great Intonations*. Ph. D. thesis, UC Santa Cruz.

Ladd, R. (2008). *Intonational Phonology*. Cambridge Studies in Linguistics.

Malamud, S. and T. Stephenson (2015). Three Ways to Avoid Commitments: Declarative Force Modifiers in the Conversational Scoreboard. *Journal of Semantics*.

Sag, I. and M. Lieberman (1975). The intonational disambiguation of indirect speech acts. *Chicago Linguistics Society*, 487–497.

Schiffrin, D. (1987). *Discourse Markers*. Cambridge: Cambridge University Press.

Stalnaker, R. (1978). Assertion. *Syntax and semantics 9*, 315–332.