Mirativity on the Table

VEDA KANAMARLAPUDI, *Stanford University*
AHMAD JABBAR, *Stanford University*

ABSTRACT

Using the Hindi-Urdu discourse particle *lo* as a case study, the present paper models mirativity within the Table model framework (Farkas & Roelofsen, 2017). This comes with an enrichment of the Table model; we propose the discourse structure to encode, in addition to discourse commitments, the time these commitments are publicized. We also incorporate a component that contains the public record of private beliefs. With these two additions, we seek to capture mirativity in the Table model, emphasizing recency as a condition for surprise and underlining the need and the conceptual difficulty of incorporating private beliefs in a public scoreboard of conversation. We show that the difficulty can be overcome to fulfill the modeling need.

1 Introduction

In this paper, we focus on *lo*, a discourse particle in Hindi-Urdu.[1] To our knowledge, *lo* hasn't been explored in the literature and hasn't even featured orthogonally in examples. We **contribute empirically** to the literature by showing that *lo* sports a rich and complex empirical profile, giving rise to a panoply of semantic puzzles, which we seek to resolve with a simple semantic generalization as in (1). We **contribute theoretically** in two ways; first, we argue that mirativity can be captured within the Table model framework (Farkas & Bruce, 2010; Farkas & Roelofsen, 2017);[2] and second, we argue that this capture requires reference to time-stamps and private belief states. Given that the components within the Table model represent the public record of the conversation, we seek to reconcile the privacy of speaker beliefs with the publicity of what the Table model seeks to capture. We introduce **Doxa**, a new Table component, and **time-stamps** to be included with each member of discourse commitments and Doxa. The reader can skip to Section 3 for the extensions.

(1) *lo* is felicitous to use at t_1, if the speaker publicly commits to p at t_1, while not having moderate to high credence towards p at t_0, where p is the content of *lo's* containing clause and $t_0 < t_1$.

We briefly foreshadow our modeling contribution for the reader's expectation of what's to be found in the coming pages in terms of theory-building. Like Gunlogson (2008), Farkas

[1]Authors contributed equally to the paper. Most special thanks to Cleo Condoravdi for her extensive feedback on this project! Thanks also to Donka Farkas, David Beaver, Miriam Butt, Rajesh Bhatt, Christine Bartels, Ashwini Deo, Seo-young Lee, Xuetong Yuan, Katherine Johnson, and audiences and participants at the Discourse Dynamics workshop at *FASAL 15* and the *Stanford Construction of Meaning* workshop! Thanks to Shrija Rawat, Azeez Hamoodur Rahman, and Mydhili Rachapudi for elaborate discussion on judgments!

[2]This point has also been made in Kraus (2019), with which we engage in Section 5.

& Bruce (2010) propose to include within the informational scoreboard (Lewis, 1979), publicized commitments of each discourse participant, in addition to what's common ground between all. Farkas & Roelofsen (2017) extend the framework in Farkas & Bruce (2010), *inter alia*, by incorporating biases in addition to full commitments, with an eye towards modeling rising declaratives and tag interrogatives.[3] One is biased towards and committed to propositions. Then, naturally, biases are represented as a set of pairs containing propositions and measures of bias, where the measure within each pair captures the extent of the bias towards the proposition in the pair. Due to wide empirical coverage, while providing a comprehensive multi-dimensional view of interpretation, Farkas & Roelofsen (2017) has emerged as a contender for a complete theory of interpretation, with an elegant division of labor between semantic composition and discourse effects.[4] To further bolster the framework's status, we argue that, with a simple enrichment, the framework can capture mirativity as well (DeLancey, 1997, 2001; Aikhenvald, 2012).[5]

What is our theoretical proposal? For each proposition p in a participant's commitment set, we propose to record the time of the participant's commitment to p. We call this time *time-stamp of commitment*, or more simply *time-stamp*. These time-stamps can slice time as thinly as one wishes; propositions might be time-stamped for the day they entered the commitment set or for the entering hour or the minute. We only make the cognitively plausible assumption that time-stamps be identified with discretized units of continuous times. We also propose the time-stamps to be recorded for biases as well. With both commitments and biases publicized at particular times, and with a public record of those times, we formalize (1) as the following: the speaker had previously not (privately) assigned moderate to high credence to the content of *lo*'s containing clause to which the speaker commits, in uttering the clause. An illustrative example will help. Suppose the speaker utters *lo S*. $[\![lo\ S]\!]$ is: the speaker had not (privately) assigned moderate to high credence towards $[\![S]\!]$ before committing to $[\![S]\!]$.[6] We also show that the optionality of privately assigning credences can be made to function coherently within the Table model, which is essentially a public record of the conversation. For this, we introduce a new component Doxa that is a public record of private beliefs. Time-stamps are extended to Doxa. Throughout the discussion, we bring Hindi-Urdu-internal and cross-linguistic evidence from English and Romance languages to motivate the extensions and understand *lo*'s contribution.

Outline of the paper. Section 2 sketches the empirical problem, while verifying (1) with each example. Section 3 proposes the extensions; sections 3.2.1 and 3.2.3 contain the time-stamps and Doxa extensions, respectively. We propose the Table-borne semantics for

[3]While Farkas & Roelofsen (2017) tie their contribution to the Inquisitive framework (Ciardelli et al., 2018), we abstract away from the Inquisitive module.

[4]The division of labor is inspired by Lauer & Condoravdi (2012) and Condoravdi & Lauer (2012) who use commitment to the content of a clause as a means for defining varying conventions of use across clause-types.

[5]Kraus (2019) attempts the same, but we seek to model mirativity more simply, while underlining the need for a reference to times and speaker private beliefs.

[6]We use $[\![\ ]\!]$ as is standard in formal semantics (Heim & Kratzer, 1998), i.e., as a function that maps expressions to their interpretations.

lo in Section 4, with some short final sections containing some reflective discussion, a comparison with Kraus (2019), and our conclusion.

2 The empirical problem

Methods. The dataset presented comprises (i) dialogues from *Dhoop Kinare* (Moin, 1987), a popular play broadcast in 1987 (with frequent reruns since) on Pakistan Television; (ii) synthetic examples with introspective judgments;[7] (iii) examples taken from the Wortschatz Leipzig Hindi Corpus of 100k sentences (Goldhahn et al., 2012; Leipzig Corpora Collection, 2019).

Design. We follow the usual strategy of constructing minimal pairs, where each pair contains examples that are identical, save for the modulation of the variable whose effect we want to measure on *lo*'s felicity. The dataset is thus structured, helping us isolate the variables that we seek to unite in an interpretable and compressed semantic generalization.

Data Distribution. Most of the attested sentences across *Dhoop Kinare* and the corpus in which *lo* occurs are declaratives. We use this statistic to motivate our focus on declaratives in sketching *lo*'s empirical profile, while also discussing interrogatives and imperatives later.

2.1 Declaratives

In (1), we presented a semantic generalization for *lo*. We motivate and verify this generalization with each context and minimal pair presented. We start with a scene from *Dhoop Kinare* (Moin, 1987, Episode 6).
DIRTY HANDS: The dad has been painting the walls distractedly, while talking to his daughter. Suddenly, the dad realizes that his hands have become all dirty from the paint.

(2) Dad: Lo, saare haath kharab ho gaye.
Dad: LO*, my hands have gotten all dirty.*

While in (2), given DIRTY HANDS , *lo* may be felicitous due to a sudden discovery, it may also just be compatible with declaratives *simpliciter*.[8] To motivate sudden discovery as a necessary condition for *lo*'s felicity in declaratives (and more broadly later), we tweak (2) a little, where the dad, instead of saying (2) upon realization, expects the hands to get dirty. This expectation is common ground before *lo* is used. Consider (3) below.

[7] We choose *Dhoop Kinare*, as it depicts the lives of young doctors in urban Karachi, with its dialogue being reflective of the Hindi/Urdu spoken across major metropolises in North India and the Punjab and Sindh provinces in Pakistan. Curiously, *Dhoop Kinare* contains many occurrences of *lo*, which many other plays we checked, even by the same playwright, couldn't match in frequency. All introspective judgments were thoroughly corroborated with native speakers of Hindi-Urdu: the authors and six other speakers from Lahore, Punjab in Pakistan and the Indian regions of Uttar Pradesh, Uttarakhand, and Hyderabad, Telangana.

[8] For instance, *lo* might be ok in declaratives by virtue of being a Force head of the sort that Bhatt & Dayal (2020) take *kya* to be. *Kya*, as per Bhatt & Dayal (2020), semantically selects for a set of propositions, but *lo*'s selection can be limited to propositions, under a framework, unlike Inquisitive Semantics (Ciardelli et al., 2018), with the type distinction between the contents of declaratives and interrogatives.

(3) Dad (before starting painting): *could you get me a towel. I always make a mess when painting*.
Daughter: *Your hands are always a sight to see*.
INTERLUDE: THE DAUGHTER GOES TO GET THE TOWEL, COMES BACK AFTER 5 MINUTES, WHILE THE DAD PAINTS DISTRACTEDLY. THE DAD LOOKS AT HIS HANDS THAT HAVE BECOME ALL DIRTY FROM THE 5 MINUTES OF PAINTING.
Dad: # Lo, saare haath kharab ho gae.
Dad: # LO, *my hands have gotten all dirty.*

Through the dialogue in (3), we seek to show that when, at time t_0, the expectation for an event is common ground, *lo*'s use, at t_1, in a speech act that proposes to make the event's occurrence common ground is not felicitous, where t_0 immediately precedes t_1. We further note that the addition of *already* makes the use of *lo* in (3) better as in (4), that may signal that the dad wasn't expecting the hands to get so dirty so quickly.

(4) Lo, abhi se hi saare haath kharab ho gae.
LO, *my hands have gotten all dirty already*.

With the examples above, we motivate a necessary condition for *lo* that the speaker commit to the asserted proposition that the speaker didn't expect to be true at the time right before. Now, we seek to assess if signaling commitment at the time of utterance is the right attitude towards the content, or if bias should suffice too.
RAIN 1: Kunal is doing his homework in the library. He decides to take a break and walk around the room. He goes to the window, and he **sees that it is raining**. Kunal says (5):

(5) Lo, bahar baarish ho rahi hai.
LO, *it's raining outside.*

To construct a minimal pair, targeting the commitment variable, contrast RAIN 1 with RAIN 2 below.
RAIN 2: Kunal is doing his homework in the library, when he decides to take a break and walk around the room. He goes to the window, and sees that it is **cloudy outside with condensation on the windowpane**.

(6) # Lo, bahar shayad baarish ho rahi hai.
LO, *it might be raining outside.*

In (2), (4), and (5), the speaker fully commits to the content of the *lo*-marked clause. In (6), there is an explicit uncertainty marker *shayad* within the *lo*-marked clause that we roughly translate in English with the epistemic modal *might*.[9] The minimal pair with (5) and (6) helps us corroborate our observation that *lo* is bad with uncertainty. In other words, *lo* requires commitment to the content of its containing clause at the time of utterance.

The above contexts, DIRTY HANDS and RAIN 1 & 2, along with their minimal pairs, suggest that commitment to p, the content of *lo*'s containing clause, at utterance time t_1, is

[9]See Jabbar (2023) for more on *shayad*, where a similar argument is built for the Hindi-Urdu discourse particle *na*.

a necessary condition for *lo*'s felicity. There's a similar question that can be raised for the immediately preceding time: at the immediately preceding time t_0, ought the speaker not be committed to p or ought the speaker not be biased towards p? Suspecting that it may be raining at t_0 is compatible with having no commitment to the proposition that it is raining at t_0. Clarifying the exact attitude whose absence *lo* requires at t_0 is thus important if we indeed see an effect of modulation of the attitude towards p at t_0 on the felicity of *lo*-p. (3), given DIRTY HANDS, is set up such that the dad has prior expectation of his hands getting dirty. This renders the use of *lo* infelicitous. Then, for the felicitous use of *lo*-p, in addition to being committed to p at t_1 (utterance time), the speaker ought not be committed to and ought not be biased towards p at t_0 (the immediately preceding time). This lack of bias and commitment towards p is exactly what makes the expression of surprise felicitous. *Lo*'s fitting this data profile can be taken as further evidence that *lo* is a mirative.

Extent of surprise. With any expression of surprise, the extent of the violation of prior expectation can be varied, which in turn can determine the felicity of the expression. For instance, in a rainy city like Manchester, if it drizzles when the forecast was for clear skies, we don't expect a life-long Mancunian to express this violation of expectation with *whoa, it is raining*. Change the scenario to a resident of the Sibi District in Pakistan; seeing rain when the forecast was for clear skies makes the use of *whoa, it is raining* quite apt.[10] Now, of course, in each case, the underlying expectation is determined through past experiences, in addition to the weather forecast. While in both cases, a prior expectation is violated, the use of *whoa* is more natural in the second instance.

What we seek to illustrate through our detour into *whoa* is that while *whoa* is an expression of surprise, the extent of the violation of prior expectation determines *whoa*'s felicity. Then, for the completeness of the semantic analysis for any mirative expression, it is important to understand if there are any conditions on the extent of surprise that warrants expression through the mirative. We set up a minimal triple below specifically to understand if there are any such conditions due to the extent of violation or the extent of surprise on *lo*'s felicity. HOMECOMING 3 is used as a control to show that in the given context, you need some expectation-violation. HOMECOMING 1 and HOMECOMING 2 seek to show that *lo* allows for a wide extent of expectation-violation.
HOMECOMING 1: A and B are roommates. A expects B to come back from work at 6pm, like usual, everyday. B comes home at 4pm today.

(7) A: Lo, tum toh 4 baje aa gaye aaj.
LO*, you're home at 4pm today*.

HOMECOMING 2: A expects B to come back from work at 4pm, but B emerges from his room at 4pm, having slept all day.

(8) A: Lo, tum toh kaam par gaye hi nahi.
LO*, you didn't even go to work today*!

[10]Sibi is considered the hotspot of Pakistan, receiving little rainfall around the year (Wikimedia Foundation, 2025).

HOMECOMING 3 (Control): A and B are roommates. A expects B to come back from work at 6pm. It is 5:55pm now. B texts A that B is nearby and will arrive in 5 minutes. B arrives after 5 minutes.

(9) A: # Lo, tum aa gae.
LO, *you're back*.

In both (7) and (8), the speaker expectation is violated. In (7), the speaker doesn't expect B to arrive early, i.e., at 4pm, and in (8), the speaker doesn't expect B to have not gone to work at all. However, arriving a few hours earlier is certainly less surprising than not having gone to work. It is felicitous to use *lo* in both situations. So, tentatively, we can say that *lo* allows for a wide range of surprise.

Now, we stop for a moment to consider two objections. The first objection goes something like this: it is not clear that the expectation-violation in HOMECOMING 1 is less severe than that in HOMECOMING 2.[11] Without a way to reliably construct contexts that vary in surprise, we cannot use the difference in felicity of *lo*, or lack thereof, across contexts to isolate a point about the range of surprise that *lo* is felicitous to use with. A different but related objection may concede that the expectation-violation in HOMECOMING 1 is less severe, but may point out that we don't exhaust the range of surprise by the two contexts, HOMECOMING 1 and HOMECOMING 2. It would then appear that the generalization that *lo* is felicitous to use in any expectation-violating context is not warranted. We consider both objections in turn.

First, to steelman the first objection, let's concede that the expectation-violation in HOMECOMING 1 is as severe as the violation in HOMECOMING 2. Even with this concession, we note that *lo* is felicitous in both the Manchester rain and the Sibi rain scenarios. This empirical fact suggests that *lo* at least accepts a wider range of expectation-violation than *whoa*. And note that here, we are not stipulating that the Manchester rain context is less surprising than Sibi rain—that would be liable to the same objection to which we are replying. Instead, we're using the variation in *whoa*'s felicity across the two contexts to take the two contexts to comprise different levels of expectation-violation. Then, we make the comparative point that *lo* allows for a wider range of surprise than *whoa*, not for its own sake, but to show that *lo* allows for a wide range of surprise. For the second objection that points out that the two contexts HOMECOMING 1 and HOMECOMING 2 don't exhaust the range of surprise, we concede that it is hard to exhaust the range of surprise through contexts. Therefore, as an approximation to its true felicity condition, we take *lo*'s felicity condition as stated below:

(10) *lo* is felicitous in contexts where the speaker is not biased towards the content of *lo*'s containing clause.

To pointedly verify (10), we construct a context where prior to committing to p, the speaker is biased towards p.

[11] Thanks to Katherine Johnson for making this point.

INFERENCE IN THE LIBRARY: Mina is doing her homework in the library, when she hears some pattering against the window.

(11) Mina: Mujhe lagta hai ki bahar baarish ho rahi hai. (looks out to check) # Lo, bahar baarish ho rahi hai.
Mina: *I think it's raining outside. (looks out to check)* # LO*, it's raining.*

As per (10), *lo* requires that the speaker not be biased towards the content of *lo*'s containing clause before utterance time. This explains (11). Now, while in INFERENCE IN THE LIBRARY, Mina was biased towards the content, the following is also quite possible. Before committing to p by way of using a *lo*-declarative, the speaker's belief state might not be sensitive to the question $?p$, where p is the content of *lo*'s containing clause. Here, in using *question-sensitivity*, we have in mind the notion explored in the philosophy and semantics literature (Yalcin, 2018; Bledin & Rawlins, 2020), whereby a belief state, construed as a set of possible worlds, might not have the resolution that makes a proposition "visible". Less abstractly, surprise can also be a reaction that stems from a general lack of awareness or due to sudden awareness about the situation.
CHRIS AT THE PARTY: A and B are going to a party. They haven't previously entertained the proposition that Chris is at the party, but they see him there. A sees Chris and says:

(12) A: Lo, Chris bhi party par hai.
A: LO*, Chris is also at the party.*

The thing to note in (12) is that A had assigned no credence (neither high nor low) to the proposition that Chris is at the party before seeing Chris. A's belief state was just not sensitive to the question *whether Chris is at the party*. So, it would be inaccurate to say that A had low prior credence to the content of (12) before uttering it. The question had simply not occurred to A. Still, the use of *lo* in (12) is felicitous. This validates (1), our generalization for *lo*, which doesn't require that the speaker have low credence towards the content of *lo*'s containing clause; instead, we framed the condition negatively, i.e., that it not be the case that the speaker had moderate to high credence towards the content at t_0. In cases like CHRIS AT THE PARTY, this condition is fulfilled, as the speaker's belief state's insensitivity to the question *whether Chris is at the party* is indeed compatible with not having assigned moderate to high credence to the proposition that Chris is at the party.

While we have looked at a wide range of contexts variable in the extent of prior speaker bias towards the content of *lo*'s containing clause, we haven't considered the scenario where the speaker has low but nonzero credence towards a proposition p before using a *lo*-declarative which denotes p.
GRADE WORRIES: A is worried about her grade in her math class, which has been quite difficult all semester long. The final scores have just come out.

(13) A: I wonder if I'll even pass the class at all … (checks her report card) Lo, main pass ho gayi!
A: *I wonder if I'll even pass the class at all …* (checks her report card) LO*, I passed!*

The speaker expresses low but nonzero credence to passing before using *lo*. We take *wonder* and *even* to be markers of low credence. It would be inaccurate to say that A has no credence assignment towards the proposition that she will pass. The proposition embedded under *wonder* is the proposition that she will pass the class, and not that she will not pass. Therefore, the expression of bias is towards the highlighted alternative (Farkas & Roelofsen, 2017) that she will pass. And we take this bias towards the proposition to be low due to its embedding under *wonder* and the use of *even*.

An interlude on credences. We pause here to note that our way of talking about credences is different from Farkas & Roelofsen (2017)'s. Farkas & Roelofsen (2017, 256) take a speaker to assign low credence to a proposition p when the speaker considers p to be somewhat more likely than the complement of p. Numerically, we understand this as still assigning >0.5 numerical credence to p. Farkas & Roelofsen (2017, 256) write "if the speaker does not consider α more likely than $\overline{\alpha}$ at all, we say that her credence in α is zero". As far as we understand this, for instances like (13) where the speaker expresses assignment of a numerical credence that is below 0.5—say 0.3 roughly—Farkas & Roelofsen (2017) would consider that numerical credence to be floored to zero on their scale.[12] While of course, this works for Farkas & Roelofsen (2017)'s purposes of understanding rising declaratives and tag interrogatives, it certainly doesn't work for *lo*. A context where the speaker assigns 0.5 to the prejacent and another context where they assign 0.3 do vary for the felicity of *lo*. So, when we say *low credence*, we have the numerical credence in mind, where we can take *low credence* to roughly lie in a range of real values between 0 and 0.4. Then, the empirical takeaway, in light of (13) is that *lo* is felicitous to use when the speaker assigns a low numerical credence to the prejacent.

So far, we've shown that *lo* goes with discovery, is incompatible with uncertainty, and is OK in both moderately unexpected and highly unexpected contexts with no prior speaker bias to the *lo*-marked proposition. We've also shown that it is OK in contexts where the speaker's belief state is not sensitive to the question of *whether p*, where p is the content of *lo*'s containing clause. We have sketched out a large part of *lo*'s empirical profile in declarative clauses, and the generalization that we have offered fits this data.

2.2 Tricky imperatives and the occasional interrogative

For completeness of our empirical analysis, we review the non-declarative examples encountered in the Wortschatz Leipzig Hindi corpus (Leipzig Corpora Collection, 2019) and *Dhoop Kinare*. First, we discuss the potential use of *lo* in imperatives. Then, we grapple with the solitary observed interrogative with *lo*.

Imperatives. Let's begin with an example where the speaker seems to use *lo* in an imperative clause.

[12] A note on em dashes: we use em dashes throughout the paper. While em dash use is often considered a sign of AI generated text (Abebe, 2025), we didn't use AI for the writing of this paper. We used AI minimally for formatting our bib entries.

DINNER: At dinner, B complains that there are only *naan*s and no *roti*s. A makes a *roti*. While handing over the *roti* to B, A says:

(14) A: Lo, yeh kha-o!
A: LO*, eat this!*

While the clause-type of A's utterance is indeed an imperative, containing the *-o* inflection on the verb that Mishra & Archana (2022) gloss as IMP.2PL, it is not clear if the *lo* used sentence-initially in (14) is indeed the discourse particle we've talked about in the context of declaratives in Section 2.1. There are a few confounding elements to note. *Lo* also occurs as the imperative form of the verb *lena*, which means *to take*. *Lena* commands require physical exchange between speaker and addressee as in the exchange of the *roti* in DINNER. We think that *lo* in (14) is not the discourse particle, but the imperative form of the verb *lena* 'to take'. We present the context DOODLE below to set up a contrast with DINNER, to strengthen our point.

In (14), we observe the *roti* being physically exchanged between speaker and addressee. Contrast this with (15), where a command is uttered by the speaker towards the addressee, but there is no exchange of a physical object.
DOODLE: A middle-schooler is doodling on an iPad. The babysitter tells them to do their homework.

(15) Babysitter: # Lo, apna kaam kar-o!
Babysitter: # LO*, do your work!*

The infelicity of (15) further corroborates that *lena* commands require physical exchange. It seems untenable to take the *lo* in (14) to be a particle use. While we don't consider examples like (14) that require an exchange of objects for understanding the distribution of *lo*, we find the following use of *lo* in an imperative environment in the Wortschatz Leipzig corpus (Goldhahn et al., 2012; Leipzig Corpora Collection, 2019) that is more congruent with the particle uses explored in Section 2.1.

(16) Lo, dekho, baarish kaun se darvaje se mere ghar mein aayi.
LO*, look from which door the rain came into my house.*

This example proves to be particularly thorny. There is no physical exchange, which we took to be the confounding factor for examples like (14).

Given that (16) is a corpus example without surrounding context, we can situate (16) in a natural context for a better understanding. The most natural context for (16) is one where the speaker of (16) did not anticipate that the rain would seep into the house through the specific door it did. The identity of the door is remarkable. (16) is felicitous only when the speaker utters it to point out the remarkability, but degraded when the remarkability comment that follows the imperative *dekho* 'see' is absent, i.e. *Lo, dekho*. Something must follow—like a remarkability comment—for the utterance to feel complete. Consider the following example, which corroborates this further.

SHOWING: A real estate agent is walking a client through a new property. The agent wants to ensure the client sees all aspects of the entryway before moving the tour along to the inner rooms.

(17) Agent: # Lo, dekhiye darvaje ko.
Agent: # LO*, look at the door.*

Lo seems incompatible with imperative clauses that aren't followed by a declarative or a rhetorical comment as in (16); when the sequence is strictly an imperative that directs the addressee to take an action as in (17), *lo* is infelicitous.[13] We invite further inquiry into the interaction between the imperatives, the remarkable comment that follows, and the felicity of *lo*; for now, we note that the felicity of *lo* in imperatives like (16) is crucially determined by whether the speaker expresses surprise or remarkability in the comment that follows the imperative. To further give a sense for the sort of environment *lo* is felicitous in, we present some English examples like (18). The reported utterance occurs in an imperative context but has the quality of pointing out something noteworthy rather than commanding an action.

DRESS: A bride-to-be goes wedding dress shopping with her entourage. Her entourage guesses that she would like to try on fitted dresses, but she comes out of the dressing room in an extravagant ballgown.

(18) Mother: Look at that gorgeous dress on you!

RECITAL: Two friends, A and B, are attending a piano recital together, where the pianist executes complex techniques during the performance.

(19) A: Listen to that *glissando*, woah.

In both (18) and (19), despite the utterance occurring in an imperative form, it serves to remark on something notable in the context. In (18), the full entourage is seated, looking at the bride-to-be enter in the ballgown. Instructing the entourage or the bride herself, literally, to *look*, would appear redundant and awkward. (19) seems felicitous to point out a remarkable execution of the *glissando* technique. The utterance in (19) seems less appropriate if literally instructing the interlocutor to listen, when the context already yields that they are doing so. This perspective is further strengthened due to cross-linguistic evidence. 2nd person imperative forms in Romance languages also contribute a mirative meaning (Villalba, 2024). Strikingly, Villalba (2024, 67) cites the marker *goita* corresponding to the verb *guaitar* 'look' in Catalan and many other miratives across Romance languages, derived from the corresponding verb in the language for 'look'. Then *dekho*/*dekhiye* 'look' is expected to have an overlapping distribution with *lo*, as the imperative verb form in the examples above can be taken to contribute a mirative meaning as well. With this provisional perspective, there really is no need to predict *lo*'s distribution in imperatives; examples like (16) are really instances of two miratives occurring together.[14]

[13] We use *direct* to include commands, requests, pleas, etc. (Kaufmann, 2011; Condoravdi & Lauer, 2012).

[14] Of course, why *lo* and *dekho* don't have a complementary distribution due to performing the same func-

Interrogatives. *Lo* appears infelicitous with interrogatives. In the Wortschatz Leipzig Hindi corpus (Leipzig Corpora Collection, 2019), there are 0 occurrences of *lo* in interrogatives. This infelicity is further corroborated by native speaker intuitions. However, we observe a single occurrence of *lo* in *Dhoop Kinare* (Moin, 1987, Episode 2):
PARKING LOT: Four friends are walking through a parking lot and find that one car in particular has been parked completely in their car's way. One of the friends, G, says (20):

(20) G: Lo ji, yeh kis-ki gaadi hai bhai?!
G: LO JI, *whose car is this* BRO?!

It is unclear whether (20) is an information-seeking question. The context is such that none of the four friends would likely be able to provide an answer to the question posed in (20), since they did not see who parked the car and do not recognize the car on sight. The question in (20) is similar to the English *Now, whose car is this?!*, which is similarly not a canonical information-seeking question. Moreover, in (20), *lo* is followed by another discourse particle *ji*. Given *lo*'s one-off and suspect distribution within interrogatives, via our semantic account, we seek to predict that it is hard to mark an interrogative with *lo*.

With a good sense of the empirical picture, which corroborates our generalization in (1), we now proceed to formalization of (1). Leveraging the argumentation above, we argue for incorporating a private beliefs component within the Table model, along with keeping track of the time when biases and commitments are made public.

3 Extending the Table model with time-stamps and Doxa

Here, we present the groundwork for formalizing the semantic generalization in (1) to precisely capture the meaning of *lo*. Our formalization assumes the discourse components in Farkas & Roelofsen (2017), who take participants, table, commitments, and evidence—all to be explained below—to be primitives. Crucially, we propose to add Doxa, a public record of private beliefs. We propose to modify commitments, evidence, and Doxa all to include time-stamps for each of their members. We formalize (1) in this enriched version.

3.1 The Table

Primitives comprise participants, table, commitments, evidence.[15]

(21) a. participants: the set of all participants in the discourse
b. table: a stack with all proposed propositions$_{inq}$[16]

tion is an interesting question, but one we don't explore currently. Jabbar & Kanamarlapudi (2023) provide evidence that particles can be stacked in Hindi-Urdu sentence-finally with propositional anaphors, so the co-occurrence of two particles is certainly not unexpected.

[15]Please note the different typeface used for the discourse components.

[16]This formulation assumes an inquisitive perspective. However, for our purposes, neither the table nor the inquisitiveness matters much. One can assume a picture where conventional update associated with declaratives happens by default, without acceptance of a proposal, and our account for *lo* should still work. Moreover, given our desire of abstracting away from the inquisitive framework for this paper, we specify inquisitive propositions with the subscript *inq* as in *propositions*$_{inq}$ to be able to more freely use *propositions*.

c. commitments: a function that takes participants in participants to sets of propositions. Where $x \in$ participants and W is the set of worlds, commitments $(x) = \{p \in \mathscr{P}(W) : x \text{ is publicly committed to } p\}$.

d. evidence: a function that takes participants in participants to sets of pairs, where each pair includes a proposition and a credence that the participant assigns to that proposition. For instance, $\langle p, [\text{high}] \rangle$ is in the set evidence (x) if x publicizes high credence to the proposition p.

Comment on (21d). We remind the reader that towards the end of Section 2.1, right after (13), we comment that our way of talking about zero, low, moderate, and high credences is different from Farkas & Roelofsen (2017)'s. We take these labels to correspond to intervals on the numerical scale from 0 to 1, whose values are used to express degrees of belief or credences in Bayesian Epistemology (Lin, 2024).

3.2 Extensions

We propose two extensions. First, we motivate the need to keep track of the time of each publicized commitment and bias. Second, we motivate the incorporation of a private beliefs component within the Table model, without which, we argue, surprise, in the way we conceive it and the speaker expresses in using *lo*, cannot be captured. We extend the time-keeping to the private belief component as well.

3.2.1 Time-stamps

(21c) and (21d) contain speaker commitments and biases that are public. As we foreshadowed above, from (21c) and (21d), it is clear that the sets to which both commitments and evidence map a participant x encode no information for *when* x publicized their commitments and biases for particular propositions. We propose to modify commitments and evidence.

(22) commitments: a function that takes participants in participants to sets of proposition-time pairs. Where $x \in$ participants and $p \in \mathscr{P}(W)$, commitments (x) $= \{(p,t) : x \text{ is publicly committed to } p, \text{ and } t \text{ is when } x \text{ publicly committed to } p\}$

(23) evidence: a function that takes participants in participants to sets of pairs of proposition-credence pairs and times. Where $x \in$ participants and $p \in \mathscr{P}(W)$, evidence (x) $= \{(\langle p,c \rangle, t) : x \text{ assigns credence } c \text{ to } p, \text{ and } t \text{ is when } \langle p,c \rangle \text{ was publicized}\}$

Explaining (22) and (23) with some illustrative examples. We co-opt an example presented by Farkas & Roelofsen (2017, 240). Consider APPLICATIONS below, and keep in mind the temporal course of the dialogue.

APPLICATIONS: Belinda and Chris are evaluating applications, when Chris expresses his exasperation by uttering (24). In an effort to keep Chris's disenchantment at bay, Belinda goes over an application and hands it over to Chris, who, thinking that the application is good, says (25). Let the dialogue be such that (24) is said at t_0 and (25) at t_1, where $t_0 < t_1$.

(24) C: None of the applications so far have impressed me.

(25) C: This is a good one ↑ ?

We can consider Chris to be publicly committed to the proposition that none of the applications have impressed him due to his utterance of (24). This can be represented as (26), where $p_1 = [\![\textit{none of the applications have impressed Chris}]\!]$. Further, we follow Farkas & Roelofsen (2017) in taking rising declaratives to signal zero to low credence.[17] Then, (25) can be represented as (27), where p_2 is the proposition that the application that Belinda has asked Chris to consider is a good one.

(26) $(p_1, t_0) \in \mathsf{commitments}(\text{Chris})$.

(27) $(\langle p_2, [\text{zero, low}]\rangle, t_1) \in \mathsf{evidence}(\text{Chris})$

The above examples illustrate how commitments and biases are to be completely represented in a system with time-stamps recording the publicizing time of the commitments and biases. Assuming a Kaplan-Lewis style two-dimensional semantics (Kaplan, 1979; Lewis, 1981), with the context c providing all the discourse components, (21a), (21b), (22), (23), and the speaker sp, the semantics for *lo* can be stated as in (28). No new typed domain is introduced for times, as we can use members of the same domain that supplies parameters to evaluate temporal expressions. However, of course, ts that are used to track the time of the publicizing of commitments and biases aren't performing the function of evaluation of propositions, but instead, that of time-stamping, as stated in (22) and (23).

(28) $[\![lo]\!]^{c,\langle w,t\rangle} = \lambda p : \underbrace{(\langle p, [\text{moderate, high}]\rangle, t_{-1}) \notin \mathsf{evidence}(sp)}_{\text{Presupposition: no moderate to high credence for } p}. \quad \underbrace{(p,t) \in \mathsf{commitments}(sp)}_{\text{Discourse effect: speaker commits to } p}$

We show that **(28) is incorrect**. We show that its failing lies in making reference to a publicized component in *lo*'s presupposition. This motivates our second extension, i.e., incorporation of a private beliefs component within the Table model.

3.2.2 The need for privacy

An important feature of both (21c) and (21d)—and (22) and (23) by extension—is that they record commitments and biases that are public. However, we think that the capture of surprise and *lo*'s distribution require reference to the private beliefs of the speaker. We argue for this from an example below. Consider the context RAIN 3 below.

RAIN 3: After a while of sitting close to the window in the living room, Gabriel decides to move to his windowless room. Michael, who is Gabriel's friend, is in Gabriel's windowless room. Gabriel knows that it is raining, something he observed when sitting by the window. However, Michael does not know this, as Gabriel hasn't publicized his commitments or

[17] As this discussion is for illustrative purposes, we stick with Farkas & Roelofsen (2017)'s way of speaking about credences, which we otherwise eschew.

bias towards the proposition that it is raining. It turns out to be infelicitous for Gabriel to reply to Michael's inquiry as in (29):[18]

(29) Michael: Is it raining outside?/I wonder whether it is raining outside.
Gabriel: # Lo, bahar barish ho rahi hai.
Gabriel: # *Lo, it is raining outside*.

However, (28) wrongly predicts *lo*'s use as in Gabriel's answer to Michael's question to be felicitous: where $p = [\![\textit{it is raining outside}]\!]$ and t_0 is utterance time, with $t_{-1} < t_0$, $(\langle p, [\text{moderate, high}]\rangle, t_{-1})$ is **not** in evidence(Gabriel), and, after Gabriel's utterance in (29), (p, t_0) **is** in commitments(Gabriel). Thus, (29) fulfills the presupposition of (28) and brings about the specified discourse effect too. But Gabriel's use of *lo* is infelicitous. The example also shows what goes wrong in general if we try to model surprise with public biases and the condition that the speaker had not previously publicized bias towards the proposition that the speaker ends up committing to by way of assertion of the proposition. To concretely make the general point, consider how it is similarly strange for Gabriel to say any of the following, in response to Michael's question, given RAIN 3.

(30) a. Wow, it is raining.
b. Whoa, it is raining.

What does the above mean for the extension to the Table model? The function evidence as laid out in (23) maps each participant to their public biases. There is no private belief component in (21)-(23). Without a reference to the private beliefs of the speaker, we can't really model speaker surprise, which should be inexpressible in cases like (29), where the speaker privately believes the proposition yet to be publicly committed to, by way asserting *lo*+p, where p = the content of *lo*'s containing clause. However, if we introduce a private belief component, we have to reconcile it with the fact that the Table model seeks to represent the public record of the conversation. As a result, we introduce Doxa.

3.2.3 Doxa

Doxa is the public record of the private beliefs of the speaker. Let's explain this further. First, note that the public record of the private beliefs of the speaker need not represent the private beliefs accurately. We propose to include Doxa as a component alongside those listed in (21).

(31) Doxa: a function that takes participants in participants to sets of pairs, where each pair includes a proposition-credence pair—where the credence encodes the participant's confidence in the proposition's truth—and a time-stamp.
If $(\langle p, [\text{moderate, high}]\rangle, t) \in \text{Doxa}(x)$, then it is public information that x is privately biased towards p at t.

[18] In this example, we have a question *is it raining outside?*. However, we have also provided an alternative: *I wonder whether it is raining outside.* Our rationale behind including the alternative is that none of the examples in earlier sections include a *lo*-declarative used as a response to a question, so the fact that *lo* is used in response to a question may be considered as a confounding factor for *lo*'s infelicity in (29).

While the members of evidence are formally similar to members of Doxa, the conditional in (31) that we isolate for pointed comment in (32) brings out the difference between evidence and Doxa.

(32) If $(\langle p, [\text{moderate, high}]\rangle, t) \in \text{Doxa}(x)$, then it is public information that x is privately biased towards p at t.

Note that in preface to (31), we noted that the public record of the private beliefs of the speaker need not represent the private beliefs accurately. Then, we also seem to take it as a sufficient condition for private bias towards p that Doxa contain p as in $(\langle p, [\text{moderate, high}]\rangle, t) \in \text{Doxa}(x)$. On the surface, this seems like an inconsistency, as the sufficiency of p's membership of $\text{Doxa}(x)$ for x's private bias towards p seems to require accurate tracking of private beliefs for Doxa. However, the consequent in (32) can be thought of as containing a publicity operator, i.e., *it is **public** information*. Then x can privately believe that the best Shakespeare play is *Measure for Measure*, while it being public information that x privately believes that *Hamlet* is the best play.

Objection. Now, an objection that can be raised right away is that if all we achieve through introduction of Doxa is the sort of private-public duality of belief as exemplified by the Bard example above, then this can also be achieved via commitments and evidence.[19] We surely don't take commitments and evidence to reliably track private states, so the duality comes by design, and expressions in language never really require reliable tracking of private belief states. Then, the objection continues: unless we show that we can indeed keep track of private beliefs—something that commitments and evidence don't record—there isn't really a point to introducing Doxa. Even if we show that such tracking is possible, to motivate Doxa, we also need to show that linguistic expressions make reference to private states. However, we aren't arguing that private beliefs can be reliably tracked. Then, what's the point of Doxa? And can Doxa be motivated on empirical grounds?

Reply. The point of Doxa is that, at utterance time, one can reveal that they privately believed or didn't believe a proposition at an earlier time—a revelation about a retrospective private belief state where the private belief was present but wasn't publicized. More explicitly: a speaker can reveal at utterance time that at an earlier time t_{-1}, the speaker didn't believe that p. Moreover, this can be done while the speaker expresses their belief that p. In fact, this is exactly what happens in many cases of mirative use; only after the speaker's utterance of mirative+p, it becomes public that the speaker didn't expect p to be true. That the speaker didn't believe p at t_{-1} was not publicized by the speaker at t_{-1}. Commitments and evidence don't contain propositions unless they are publicized. One can further resist introduction of Doxa on the grounds that we can allow commitments and evidence to contain earlier times too. However, $(\langle p, [\text{moderate, high}]\rangle, t_{-1}) \in \text{evidence}(x)$ would simply mean that x publicized their bias towards p at t_{-1}. Commitments and evidence, by design, are such that things cannot enter them paired with t without getting

[19] There's a lot of work that points out the public-private duality of belief (Yalcin, 2012; Goldstein & Kirk-Giannini, 2022; Bary, 2025), often to ask the reader not to confuse common ground with common belief.

publicized at t. Whereas, one need not publicize that they believe that p for it to later enter Doxa that they privately believed that p at an earlier time. One publicizes the commitment to or bias towards propositions that go into commitments and evidence, but one publicizes information that reveals that one is/was privately biased towards the proposition and this proposition goes into Doxa—not only with the time of the utterance, but possibly also an earlier time. This shows conceptually that Doxa captures something that commitments and evidence don't. We bring out these differences empirically in the calculations in Section 4. Now, we move to the unique empirical coverage that Doxa brings.

Unique empirical coverage. We want to be able to say that in (29), by using *lo* as in *lo, it is raining*, Gabriel reveals to Michael, i.e., publicizes, that Gabriel didn't privately believe right before utterance time that it is raining. This is what contributes to *lo*'s infelicity, as Gabriel is in a windowless room and it is not possible for him to form the belief that it is raining, on the spot. The same point can be made with *whoa* and *wow*—at least for examples like (30)—where the speaker publicizes that they didn't privately believe the prejacent right before committing to the prejacent. Of course, in our discussion here, we have used *believe*, instead of talking about biases or credences, which is exactly what we shift to, in our explanations and calculations in Section 4.

To conclude this section, we emphasize that it is the public recording of whether the speaker possessed a certain belief privately at a time earlier than utterance time that the components in (21) cannot do. Moreover, it is quite compatible with the aims of the Table model, i.e., modeling conventional discourse effects, to make reference to private beliefs in a discourse component like Doxa if an expression can be taken to conventionally encode the information that at a prior time the speaker didn't hold a certain belief privately. This information is made public via the use of the expression.

4 The Table-borne semantics for *lo*

At the outset in Section 1, we stated the semantic generalization for *lo*. We repeat it below.

(33) *lo* is felicitous to use at t_1, if the speaker commits to p at t_1, while not having moderate to high credence towards p at t_0, where p is the content of *lo's* containing clause and $t_0 < t_1$.

Through our discussion in Section 3.2.3, we argued that it need not be public that the speaker does not have moderate to high credence towards p at t_0; instead, **this lack of bias could be private until it is revealed at utterance time**. The crucial thing to note is that with the use of *lo*, the speaker signals that they weren't biased towards p before they committed to it. This past lack of bias can either be signaled by having explicitly expressed bias towards the complement of the proposition or be signaled at utterance time by the use of *lo* itself. The public record that the speaker doesn't have the private belief that p at t entails that the speaker does not have any public bias towards p at t. We can lay this out more explicitly in (34) and (35). Where c is a Kaplanian context and t is a time s.t. t may be a time other than the one initialized by c:

(34) $\forall c \forall t : p \in \mathsf{evidence}_{c,t}(x) \rightarrow p \in \mathsf{Doxa}_{c,t}(x)$

(35) $\forall c \forall t : p \in \mathsf{commitments}_{c,t}(x) \rightarrow p \in \mathsf{Doxa}_{c,t}(x)$

With the above specifications in place, we can say that the sub-condition "while not having moderate to high credence towards p at t_0" in (33) should be modeled with Doxa.

We have verified the generalization in (33) with all the data presented in Section 2. We present a formalization of (33) now. We assume a two-dimensional semantics in the Kaplan-Lewis style (Kaplan, 1979; Lewis, 1981). We take c to denote the context of utterance, which includes, *inter alia*, the speaker of the utterance, notated as sp, and all the primitives of our version of the Table model, i.e., participants, table, commitments as in (22), evidence as in (23), and crucially Doxa as in (31).

(36) $[\![lo]\!]^{c,\langle w,t\rangle} = \lambda p : \underbrace{(\langle p, [\text{moderate, high}]\rangle, t_{-1}) \notin \mathsf{Doxa}(sp)}_{\text{Presupposition: no moderate to high credence for } p}. \quad \underbrace{(p,t) \in \mathsf{commitments}(sp)}_{\text{Discourse effect: speaker commits to } p}$

We show the workings of (36) with a few calculations. Let's first start with (29)—the problematic case that motivated the incorporation of the private beliefs component.

(37) $\lambda p : (\langle p, [\text{moderate, high}]\rangle, t_{-1}) \notin \mathsf{Doxa}(sp).(p,t) \in \mathsf{commitments}(sp)(\textit{it is raining})$
Discourse effect: $(\textit{it is raining}, t) \in \mathsf{commitments}(sp)$
Presupposition: $(\langle \textit{it is raining}, [\text{moderate, high}]\rangle, t_{-1}) \notin \mathsf{Doxa}(sp)$

(36) shows exactly why Gabriel's use of *lo* is infelicitous in (29). We specified in (29) that Gabriel moves into a windowless room where Michael asks him if it is raining. There is no obvious way—at least obvious to Michael—for Gabriel to *not* have privately believed, right before utterance time, that it is raining in order to be able to cooperatively assert that it is raining at utterance time.[20] That is what makes the accommodation of the Presupposition in (37) hard too, leading the addressee to think that Gabriel knew that it was raining before utterance time too. If, for instance, Gabriel instead prefaces his answer with *let me check the weather app* and finds out that it is raining, upon formation of this belief, Gabriel can indeed use *lo*, and in that instance, Gabriel's utterance as in (29) would be felicitous.

Consider another example. In (11), the speaker, Mina, publicizes her bias towards the proposition that it is raining by saying *I think it's raining outside*. The proposition that it is raining, paired with high credence, enters evidence(Mina), and by the entailment outlined in (34), Doxa(Mina) as well. As this expression of bias occurs before the utterence of *lo*'s containing clause, the presupposition $(\langle \textit{it is raining}, [\text{moderate, high}]\rangle, t_{-1}) \notin \mathsf{Doxa}(\textit{Mina})$ is not met, predicting infelicity for the use of *lo* in a following utterance as in *Lo, it is raining outside*. And indeed this utterance in (11) is infelicitous.

We also use (36) to show why *lo* doesn't go well with interrogatives. Take a canonical information-seeking question such as *is it raining?*. (36)'s discourse effect is that the speaker commits to the prejacent, which we can take to be the proposition denoted by the

[20] By *cooperative* as in *cooperatively assert*, we specify assertion of propositions for which the speaker has adequate evidence (Grice, 1975).

sentence radical *is it raining?*.[21] Then, in requiring commitment to the sentence radical, *lo*'s use in an interrogative, whose purpose is to seek information, turns out to be infelicitous. The examples considered above show the role of Doxa, how it interacts with evidence as specified in (34), and how in making reference to Doxa in the semantic clause for *lo* as in (36), we capture data like (29), which motivated the need for incorporation of a private beliefs component. As our contribution seeks to model mirativity within the Table framework, we compare our account with Kraus (2019), who seeks to do the same.

5 Comparison with Kraus (2019)

Kraus (2019)'s contribution. Kraus considers the cases of English miratives *oh* and *huh*. Kraus argues that these miratives require speaker's expectation-violation. Where p is the content of the mirative's containing clause and $\mathsf{Exp}(p)$ is supposed to capture speaker's expectation for p, Kraus proposes the mirative to contribute the following change to the state of the discourse: $\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)$ is added to the speaker's discourse commitments, i.e., the changed discourse contains $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)] \in \mathsf{commitments}(sp)$ as the new information. Further, Kraus thoroughly considers a range of prosodic data, while also nicely capturing a subtle difference between the contribution of *oh* and *huh*. For this, Kraus relies on the divergence in examples such as (38)—taken from Kraus (2019, 26)—to construct a slightly modified effect for the two, where in felicitous uses of *huh*—but not of *oh*—, the speaker must have low, but nonzero credence towards the surprising proposition, i.e., $\mathsf{Exp}(p) > 0$.

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

Problems. First, we note that Commitments for a participant x is intensionally defined as containing the propositions that x is publicly committed to.[22] And $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$ is not a proposition. There's another worry here. Kraus's modification to the Table includes a stack of pending at-issue propositions yet to be accepted or rejected. Further, Kraus requires that $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$ continue to remain in $\mathsf{commitments}(sp)$ even after p is resolved, i.e., taken off the Table. Then, if $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$ seeks to model that the speaker doesn't expect p to be true, then $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$'s persistence in $\mathsf{commitments}(sp)$ creates a discourse context where p is resolved and the speaker still doesn't

[21] Roelofsen & Van Gool (2010), and later Farkas (2011), Roelofsen & Farkas (2015), and Farkas & Roelofsen (2017) appeal to a level of semantic representation that they call *highlighted content*. Theiler (2021) makes a compelling case for making reference to highlighted content for explaining the semantic contribution of the discourse particle *denn* and *überhaupt* in German. Jabbar & Kanamarlapudi (2024) do the same for the Hindi-Urdu *bhala*. While we opt for the simpler option currently with sentence radicals, one can, in principle, use highlighted propositions in (36) and the calculations.

[22] We would like to acknowledge that the problems raised in these paragraphs are based on our read of Kraus (2019), which can very well be incorrect.

expect p. Not only is this context state undesirable on its own, it leads to the prediction that the speaker could in principle keep expressing surprise felicitously at repeated assertions of the proposition at which the speaker expressed surprise using *oh* at first. To see how this follows, consider (39) from Kraus (2019, 29).

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

Kraus (2019, 30) proposes that, after B's utterance of $oh{+}q$, commitments(B) contains: q, $\mathsf{Exp}(q) \approx 1$, and $[\mathsf{Exp}_B(p) < \mathsf{Exp}_B(\neg p)]$. Kraus specifies that the common ground is unchanged. p that was put on the table after A's assertion is popped after B's utterance of $oh{+}q$. Then, later in the conversation, if A were to assert p again, given that $[\mathsf{Exp}_B(p) < \mathsf{Exp}_B(\neg p)]$ is still in commitments(B), B can express surprise again with *oh*.

Our account differs from Kraus's also in that we consider miratives in environments with discovery of a fact, whose propositional content p may not be at-issue, and instead the speaker uses the mirative with p as in $lo{+}p$ or $oh{+}p$. Consider (7), where the speaker says *Lo, you're home at 4pm today*, where the proposition that the addressee is home at 4pm today is surprising. All of Kraus's example include contexts where the speaker of the mirative uses the mirative to settle the issue at hand, expressing surprise at the content of the issue on the table, and not at the content of the mirative's containing clause. Our account doesn't require the restriction that the speaker of a *lo* utterance be responding only to an utterance relayed to them by an interlocutor.

Given that in (7), the speaker can say *oh, you're back home at 4pm today*, Kraus's account makes the following prediction, we think. Where $p = [\![$*you're back home at 4pm today*$]\!]$, the speaker utterance *oh, you're back home at 4pm today* not only adds $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$ to speaker's commitment set, but also the content of *oh*'s containing clause, i.e., p. This leads to a commitment set with both $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$ and p, meaning that the speaker doesn't expect p and that the speaker is committed to p. This is surely to be avoided.

Another issue due to $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]$'s inclusion in the speaker's commitment set due to using a mirative is that there's a considerable chunk of contexts like (12) that Kraus rules out as bad for mirative use, which in fact license the use of miratives. These are contexts where the speaker may not have any attitude towards the event or proposition at which the speaker expresses surprise using the English *oh* or the Hindi-Urdu *lo*. Such contexts may be taken to be instances of *sudden awareness* (Aikhenvald, 2012, 462), or more formally, as those where the speaker's belief state is just not sensitive to the relevant question (Yalcin, 2018; Bledin & Rawlins, 2020). In such contexts, the speaker can use the English *oh* and the Hindi-Urdu *lo*. The characteristic feature of examples like (12) is that the speaker had no bias either for p or $\neg p$, when the speaker ended up using $lo{+}p$.

Remedy. Currently ignoring that commitments require that its members have a certain type, we think that Kraus's proposal can be made to work, for the most part, if we include time-stamps.[23] More specifically, if by using a mirative, the speaker expresses $[\mathsf{Exp}(p) <$

[23]We add the qualification *for the most part*, as our suggested remedies here don't account for all of the raised objections.

$\mathsf{Exp}(\neg p)]_{t_{-1}}$, where the subscripted t_{-1} indicates that this expectation held at t_{-1}, a time before utterance, then the speaker commitments set can coherently contain both $[\mathsf{Exp}(p) < \mathsf{Exp}(\neg p)]_{t_{-1}}$ and p which was publicized at utterance time, as long as commitments records p with its publicizing time, i.e., t. We propose exactly this innovation in our paper. We think that, while Kraus doesn't make reference to time-stamps in the modeling, Kraus is presciently aware of their importance; Kraus (2019, 25) writes: "*Oh* indicates that the speaker **had** higher expectations for $\neg p$ over p" (emphasis added). The use of the past tense, even when not modeling it formally, further underscores the need for making reference to the time when each member of commitments, evidence, and Doxa entered these sets.

6 Conclusion

In addition to mapping out the empirical profile of a previously unexplored discourse particle in Hindi-Urdu, we have argued for the necessity of time-stamps for modeling mirativity within the Table model. We brought this out further in engagement with Kraus (2019), who seeks to model mirativity without reference to times. We highlight a few issues with Kraus's approach, many of which, we show to be resolvable by the simple inclusion of time-stamps. Further, with our discussion in Section 3.2.2 of examples like (29), we showed that speaker's previously unpublicized beliefs can be revealed—interestingly when the speaker no longer holds these beliefs. In this way, the speaker never publicizes these beliefs, but they come to be on public record, as constitutive of the speaker's belief state at an earlier time. Again, it is time-stamps that afford us the expressive power to be able to say the above. We hope that our contribution invites the semantics community to make further progress on modeling mirative expressions within the Table model. We also think that our inclusion of time-stamps and the public record of private beliefs via Doxa provides an opportunity to explore the meta-theory of the Table in future work.

References

Abebe, Nitsuh. 2025. With the Em dash, A.I. embraces a fading tradition. *The New York Times Magazine* `https://www.nytimes.com/2025/09/18/magazine/chatgpt-dash-hyphen-writing-communication.html`.

Aikhenvald, Alexandra Y. 2012. The essence of mirativity. *Linguistic typology* 16(3). 435–485.

Bary, Corien. 2025. Speech acts, common ground and commitments. *Linguistics and Philosophy* 48(3). 505–526.

Bhatt, Rajesh & Veneeta Dayal. 2020. Polar question particles: Hindi-urdu *kya*. *Natural Language & Linguistic Theory* 38. 1115–1144.

Bledin, Justin & Kyle Rawlins. 2020. Resistance and resolution: attentional dynamics in discourse. *Journal of Semantics* 37(1). 43–82.

Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen. 2018. *Inquisitive semantics*. Oxford University Press.

Condoravdi, Cleo & Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. *Empirical issues in syntax and semantics* 9. 37–58.

DeLancey, Scott. 1997. Mirativity: The grammatical marking of unexpected information. *Linguistic Typology* 1. 33–52.

DeLancey, Scott. 2001. The mirative and evidentiality. *Journal of Pragmatics* 33(3). 369–382.

Farkas, Donka F. 2011. Polarity particles in English and Romanian. In *Romance linguistics 2010*, 303–328. John Benjamins Publishing Company.

Farkas, Donka F & Kim B Bruce. 2010. On reacting to assertions and polar questions. *Journal of Semantics* 27(1). 81–118.

Farkas, Donka F & Floris Roelofsen. 2017. Division of labor in the interpretation of declaratives and interrogatives. *Journal of Semantics* 34(2). 237–289.

Goldhahn, Dirk, Thomas Eckart, Uwe Quasthoff et al. 2012. Building large monolingual dictionaries at the Leipzig corpora collection: From 100 to 200 languages. In *LREC*, vol. 29, 31–43.

Goldstein, Simon & Cameron Domenico Kirk-Giannini. 2022. Contextology. *Philosophical Studies* 179(11). 3187–3209.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and semantics, vol. 3: Speech acts*, 41–58. New York: Academic Press.

Gunlogson, Christine. 2008. A question of commitment. *Belgian Journal of Linguistics* 22(1). 101–136.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Malden, MA: Wiley Blackwell.

Jabbar, Ahmad. 2023. The Hindi-Urdu *na* and reasonable inference. *Proceedings of the 59th Chicago Linguistic Society (CLS 59)* Unpublished manuscript.

Jabbar, Ahmad & Veda Kanamarlapudi. 2023. Grounding with particles. In *Proceedings of the 27th workshop on the Semantics and Pragmatics of Dialogue (SemDial 27)*, .

Jabbar, Ahmad & Veda Kanamarlapudi. 2024. *Bhala*: Subordinating highlighted propositions. *Presentation at the 60th annual meeting of the Chicago Linguistic Society (CLS 60)* .

Kaplan, David. 1979. On the logic of demonstratives. *Journal of Philosophical Logic* 8. 81–98.

Kaufmann, Magdalena. 2011. *Interpreting imperatives*, vol. 88. Springer Science & Business Media.

Kraus, Kelsey. 2019. Intonation and expectation: English mirative contours and particles. In *Proceedings of sinn und bedeutung*, vol. 23 2, 19–36.

Lauer, Sven & Cleo Condoravdi. 2012. The basic dynamic effect of interrogative utterances. In *13th texas linguistics society conference, university of texas, austin*, .

Leipzig Corpora Collection. 2019. Hindi web text corpus (india), 2019. `https://dict.wortschatz-leipzig.de?corpusId=hin-in_web_2019`.

Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8. 339–359.

Lewis, Davis. 1981. Index, context, and content. *Philosophy and grammar: Papers on the occasion of the quincentennial of Uppsala University* 79–100.

Lin, Hanti. 2024. Bayesian Epistemology. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford encyclopedia of philosophy*, Metaphysics Research Lab, Stanford University Summer 2024 edn.

Mishra, Anushree & S Archana. 2022. An analysis of imperatives in Hindi-Urdu. *Coyote Papers* 60–71.

Moin, Haseena. 1987. Dhoop Kinare. Television drama series. Screenplay by Haseena Moin. Broadcast on Pakistan Television (PTV).

Roelofsen, Floris & Donka F Farkas. 2015. Polarity particle responses as a window onto the interpretation of questions and assertions. *Language* 91(2). 359–414.

Roelofsen, Floris & Sam Van Gool. 2010. Disjunctive questions, intonation, and highlighting. In *Logic, language and meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, revised selected papers*, 384–394. Springer.

Theiler, Nadine. 2021. *Denn* as a highlighting-sensitive particle. *Linguistics and Philosophy* 44(2). 323–362.

Villalba, Xavier. 2024. Romance exclamative markers at the syntax-pragmatics interface: A compositional approach to exclamativity. *Journal of Pragmatics* 226. 64–77.

Wikimedia Foundation. 2025. Sibi district. `https://en.wikipedia.org/wiki/Sibi_District`. Accessed: 2025-11-09.

Yalcin, Seth. 2012. Bayesian expressivism. *Proceedings of the Aristotelian Society* 112(2). 123–160.

Yalcin, Seth. 2018. Belief as question-sensitive. *Philosophy and Phenomenological Research* 97(1). 23–47.