

Comma intonation as speech-event indefiniteness

word count: 9,116

Abstract

The goal of this paper is a non-stipulative semantic account of parentheticals. In addition to accounting for observed syntactic and semantic properties of parentheticals, we want to explain the fact that parentheticals are obligatorily offset by pauses (i.e. comma intonation), and that this mechanism for marking parentheticality seems to be cross-linguistically universal. In the past, researchers have avoided ascribing too much semantic significance to pauses because not all pauses form parentheticals; we solve this problem by analyzing pauses as speech-event indefinites (following recent formal accounts of speech events; Bittner, 2007; Rett, 2024). Any pause introduces a new speech event discourse referent that may or may not be coextensional to a previous speech event. From this perspective, we can simultaneously account for self-interruption and many more non-construction-specific instances of pauses as well as parentheticals and appositives.

1 Introduction

The goal of this paper is a compositional semantic account of parentheticals, in particular one that avoids stipulating null operators, and one which takes into account the (obligatory) role of comma intonation, i.e. pauses, in their formation.

We use the term ‘parenthetical’ to characterize any phrase or clause offset parenthetically. Orthographically, this parenthetical status is typically marked with parentheses or commas; in speech, it is signified by comma intonation. A sentence-medial parenthetical is offset prosodically with pauses on either end.¹ (1) exemplifies a sentence-medial parenthetical, in particular, an appositive relative clause.

- (1) Andi said that a Canadian, who was talking about Kant all night, was the only philosopher at the party.

Parentheticals typically contribute not-at-issue (NAI) content, which is to say that in standard contexts (1) makes the same assertion that it would without the parenthetical. With or without the parenthetical, (1) reports a speech event of Andi claiming a Canadian was the only philosopher at the party. The additional information encoded in the parenthetical (that the Canadian was talking about Kant all night) does not have the ability to falsify the claim about what Andi said, and intuitively functions as a side comment reflecting the speaker’s knowledge about the referent rather than as part of what is attributed to Andi.

In contrast, (2) is a non-parenthetical version of (1). It differs only in that it does not include any comma intonation offsetting the relative clause.

¹We are, of course, not in the business of assigning semantic importance to orthography; what’s crucial to us is the prosodic pauses. We use the term ‘comma intonation’ in deference to prior literature, but it should be taken with a healthy grain of salt.

- (2) Andi said that a Canadian who was talking about Kant all night was the only philosopher at the party.

This is enough to change the status of the relative-clause content to at-issue (AI). In other words, without the pauses off-setting the relative clause, the relative clause content is part of the report of Andi’s speech, and part of the assertion. Whether or not Andi indeed mentioned that the Canadian was talking about Kant all night now has the ability to falsify the sentence.

Because of their non-restrictive character, appositive relative clauses can occur on individual-denoting expressions that don’t ordinarily combine with relative clauses (though see Sharf 2024a,b):

- (3) a. Andi said that Adam, who was talking about Kant all night, was the only philosopher at the party.
 b. *Andi said that *(the) Adam who was talking about Kant all night was the only philosopher at the party.

In fact, virtually any kind of expression can be offset by comma intonation, and therefore virtually any kind of expression can form a parenthetical. Some additional examples are in (4).

- (4) a. The Canadian talking about Kant all night was, frankly, not invited. *discourse adverb*
 b. The Canadian, the only philosopher at the party, was talking about Kant all night. *nominal*
 c. The Canadian was talking about Kant all night, as philosophers often do. *as-clause*
 d. The Canadian, if that is indeed where he is from, was talking about Kant all night. *if-clause*
 e. The Canadian (I was born in Canada too) was talking about Kant all night. *full clause*

We want to add to this list cases of self-interruption, because they also require comma intonation:

- (5) The Canadian—sorry, it’s so hot in here, I’m sweating—was talking about Kant all night.

Like the examples in (1) and (4), the self-interruption in (5) is a side comment, i.e. intuitively not part of what the utterance asserts. And it is, of course, (necessarily) offset by comma intonation.

Cases like (5) have traditionally not been considered parentheticals because their content doesn’t directly bear on any other part of the meaning of the sentence in which they occur. Intuitively, the difference between a self-interruption and a parenthetical is that the former introduces a separate issue. But our characterization of parentheticality doesn’t mention relevance or topicality, and we’ve found none that explicitly do.² So, in the absence of any evidence that languages employ distinct strategies for forming canonical parentheticals than they do for self-interruption, (5) is a natural extension of the core phenomenon, and an ideal theory is one that attributes to the common mechanism a common effect.

As we will explain in detail in §2, there are roughly three classes of parenthetical accounts on the market. The first originate largely in the intonation literature, and focus on the role of comma intonation in the construction of parentheticals (Selkirk, 1995; Truckenbrodt, 2015). These accounts consider comma intonation to demarcate separate speech acts, but do not define what it means for a single utterance to contribute more than one speech act, nor do they provide a compositional semantic treatment of parentheticals. The second class, exemplified by Koev (2012, 2013), do proffer formal semantic accounts of parentheticals, but ignore the role of comma intonation, and are thereby forced to cleave the natural class of data in (4) into ‘pure’ and ‘impure’ parentheticals.

²It is hard to imagine what such a definition of parentheticality would look like, given that relevance and topicality is gradient: whether or not one semantic content pertains to another semantic content is not a binary issue.

We count ourselves among a third class of accounts, which focus on the role of comma intonation in the formation of parentheticals and aim to provide a compositional account for them. This class also includes the account of comma intonation in AnderBois et al. (2010, 2013), wherein left commas signal a switch from at-issue to not-at-issue content, and right commas signal a switch back. As we’ll argue in §2.2.1, this doesn’t result in empirically accurate predictions.

The present proposal blends the insights of the intonation-based speech act accounts of parentheticals with recent dynamic-semantic accounts of discourse anaphora to the present speech event (Bittner, 2007; Rett, 2024). The insight is that, just as interlocutors keep track of discourse participants when they use first- and second-person anaphors, they also keep track of different speech events (and their content, speaker, and runtime). We propose that speakers across languages use pauses to demarcate speech events, and that any pause—including those between sentences—introduces a distinct speech event discourse referent (‘dref’) that can potentially refer to a new speech event. Parentheticality is just what happens when the speech event following a pause is interpreted as a distinct event from the one that preceded it, rather than a continuation of the same event.

The important difference between a speech act and a speech event is that there is evidence that the latter is linguistically modeled and restricted in discourse anaphora and other phenomena (Rett, 2024). We argue that, in addition to being a truly compositional semantic treatment of all parentheticals, this approach can account for the at-issue flexibility of parentheticals, the phenomenon of self-interruption, and the universality of the use of comma intonation to demarcate parentheticality. It is also incredibly versatile: we can compositionally extend this account seamlessly to any pause.

The paper proceeds as follows: in §2 we briefly review prior semantic accounts of parentheticals, and argue that they are either formally or empirically insufficient. Then in §3 we present our account, starting with a review of formal analyses of the speech event. In §4, we describe the empirical reach of the account, measuring our theory by the desiderata set out in Koev (2013), and also talking about cases of self-interruption. §5 concludes.

2 Prior accounts

2.1 Prosodic accounts

We focus on two papers in the prosody literature that take up parentheticals quite seriously: Selkirk (1995) and Truckenbrodt (2015). Both papers focus on the intonational phrase, and argue that each intonational phrase corresponds to a different speech act. They both do so without really providing an account of what qualifies as a speech act, and how a single utterance can contain more than one (although see Searle and Vanderveken, 1985, for such a proposal).

An intonational phrase is a prosodic unit offset by significant pauses on either side, and bearing an instance of nuclear stress anchoring a terminal contour (see Jun 2022 for a useful overview). Typically only one intonational phrase is required to prosodically parse a sentence, though one can optionally be parsed into multiple intonational phrases. Selkirk and Truckenbrodt argue that parentheticals are clearly parsed into a distinct intonational phrase from the main clause, separated as they are from the main clause by significant pauses and bearing their own nuclear stress.

While Selkirk and Truckenbrodt both rely on an intuitive understanding of what a speech act is, Truckenbrodt offers several related but compelling tests for the illocutionary independence of parentheticals. His tests largely involve the distribution of incompatible illocutionary markers in German (e.g. *ja* and *wohl*); Truckenbrodt assumes these discourse particles contribute meaning at the speech-act level. Because they encode contradictory meaning—*ja* marks speaker certainty, while *wohl* encodes a weaker epistemic relation between the speaker and the propositional

content—they cannot cooccur in the same sentence. Truckenbrodt (2015) notes that parentheticals are an exception to this rule, as in (6).

- (6) Peter, der **ja** gerade in Berlin ist, hat **wohl** schon angerufen.
 P who PRT currently in B is has PRT already called
 ‘Peter, who *as we know* is in Berlin, has *I supposed* already called.’ *German*
 (Truckenbrodt, 2015)

Truckenbrodt thus interprets the acceptability (and coherence) of (6) as evidence that the parenthetical doesn’t just introduce a distinct proposition, but that it does so by introducing a distinct speech act (presumably, of assertion), perhaps equivalent to a succession of speech acts, as in (7).

- (7) I suppose Peter has already called.
 We know Peter is in Berlin.

Though our account will pursue the same intuition as Selkirk’s and Truckenbrodt’s, there are several problems with this approach as they articulate it. First, the notion of a speech act is unfortunately quite formally messy, and it’s unclear how to implement the introduction of a speech act compositionally without stipulating otherwise unmotivated null operators. (For an overview of problems with this approach, see Extepare, 1997).

Second, there is nothing in Selkirk’s or Truckenbrodt’s account that can differentiate between the potential (not-)at-issue status differential between a main clause and a parenthetical one: the two assertions in (7) are equal in their at-issue status, and differ only in the linear introduction of each assertion.

Finally, in Selkirk’s and Truckenbrodt’s accounts, the relationship between parentheticality and intonational phrasing is unidirectional: parentheticals must be parsed into their own intonational phrase, but breaking up a sentence into multiple intonational phrases doesn’t necessarily induce parentheticality. Speakers can pause mid-sentence at essentially any time, and for essentially any reason. Take the following case for example:

- (8) a. Nobody is both a good linguist and a good philosopher
 b. Nobody is both a good LINGUIST, and a good PHILOSOPHER
 c. NOBODY, is BOTH, a good LINGUIST, AND, a good PHILOSOPHER

This sentence can be parsed into a single intonational phrase (8a); or, perhaps for emphasis, broken up into two (8b); or, perhaps excessively theatrically, into five (8c), and so on. It’s not how we typically speak, but it’s a completely grammatical and interpretable way of speaking. But, for all of the sentences in (8), it’s clear that a single assertion of a single proposition is involved; compositional semantic processes like the scope of negative quantification over the conjunction are preserved across the intonational phrases even in the most extremely subdivided pronunciation of the sentence.

Selkirk and Truckenbrodt were well aware of this issue, and realized that they could not say that *any* pause delineates a distinct speech act. Truckenbrodt is particularly clear in stipulating an interpretational difference between ‘optional’ and ‘obligatory’ pauses: obligatory pauses demarcate a new speech act, while optional pauses are semantically vacuous.

But while it seems plausible that a language user knows, for any given pause, whether or not it was obligatory (i.e. whether or not it contributed to a parenthetical), it’s not clear how we would represent that knowledge compositionally. We suggest that Truckenbrodt’s approach gets things the wrong way around: rather than analyzing parentheticality as the interpretation of ‘obligatory pauses’, we seek a general theory of pauses that derives the prediction that parenthetical inter-

pretations *require* pauses, while pauses do not *necessitate* parentheticality. The present proposal builds on the insights of these accounts while avoiding the messy notion of a speech act, and while managing to treat all pauses as semantically identical.

2.2 Semantic accounts

2.2.1 Conventionally NAI accounts

Famously, Potts (2005) added meat to the bones of the Gricean idea that conventional implicature is distinct from other conventionally-encoded content. In the class of conventional implicature, Potts grouped parentheticals in with expressives like *bastard* and multi-dimensional operators like *but*. Potts hard-wired the not-at-issue (NAI) status of all of these elements into his multi-dimensional semantics, and the resulting theory has built into its core the prediction that all parentheticals encode NAI content.

The proposal in Murray (2010, 2014) makes the same prediction, but in a different theoretical framework. Murray proposes a formal, dynamic account of utterances and utterance types based on the manipulation of the Stalnakerian Common Ground (CG): not-at-issue content directly updates the CG; at-issue assertions amount to a proposal to update the CG; and interrogatives partition the CG. In this framework, as in Potts (2005), parentheticals are treated as just another kind of conventionally NAI content, and as such, they directly update the CG. Also as in Potts (2005), there is no discussion of how parentheticals come to acquire this NAI status, so both proposals additionally fall short of being meaningfully compositional.

The account of parentheticals in AnderBois et al. (2010, 2013) is similar to Murray’s in that it works within the framework of dynamic Stalnakerian update, and in that it hard-wires NAI status into parentheticals. But it differs in that the authors identify comma intonation as the compositional piece of the parenthetical puzzle. For them, commas toggle back and forth between AI and NAI content (see also Nouwen, 2007), like (other) “propositional operators”, e.g. negation and modals: “both appositives and propositional operators are ultimately taken to *toggle* between the main proposal and some other propositional dref. Where appositives are special is that they toggle to the dref for the Context Set itself (i.e., p^{cs}) rather than introducing a new dref p' ” (p26). In order for this to work, they must assume that left commas and right commas are distinct, in that the former goes from AI $\rightarrow$ NAI, and the latter from NAI $\rightarrow$ AI. “That is, a left comma operator indicates that the immediately following content should update p^{cs} , while a right comma operator toggles back to updating p itself” (AnderBois et al., 2013, 17). They do not, however, explicitly define either.

These accounts all have in common the foundational assumption—and, by design, the shared prediction—that any given parenthetical necessarily introduces NAI content. But we now know, from a line of research beginning with Tonhauser et al. (2013), that very few meanings are conventionally backgrounded: at-issue status is generally flexible, and can come and go with respect to the Question Under Discussion.

And we now also know that this also holds for parentheticals specifically. Wang et al. (2005) point out that the nominal parenthetical in (9) is more naturally interpreted as part of the belief report than as a side comment reflecting the speaker’s perspective (in contrast with (1)).

(9) John believes that a professor, a quite famous one, published a new book.

And, as extensively discussed in Esipova (2018), parentheticals can in fact do at-issue things: they can be the target of polarity particle anaphora (10), or they can answer QUDs under the right circumstances (11):

- (10) a. A: Paula, who brought cookies, left after midnight.
 B: ?No she didn't, she brought wine.
 b. A: I met Paula, who brought cookies.
 B: No she didn't, she brought wine.
- (11) a. A: What did Paula bring and when did she leave?
 B: ?Paula, who brought cookies, left after midnight.
 b. A: Who did you meet at the party and what did they bring?
 B: I met Paula, who brought cookies.

In (10a), the polarity particle in B's response can't target the parenthetical content, which is what we expect of NAI content. But this changes in (10b), where the polarity particle in B's response can successfully target the parenthetical content. The clear difference between the two examples is the location of the parenthetical; the felicity of (10b) suggests that the issue with (10a) is that the parenthetical is not the most salient target for propositional anaphora due to its non-final position, rather than supporting the conclusion that parentheticals are intrinsically NAI.

Similarly, NAI content is unable to address the QUD. The infelicity of (11a) might be taken to demonstrate that parentheticals encode NAI content. But again, things change when the parenthetical is sentence-final; the response in (11b) is totally acceptable. This suggests that the problem with (11a) is that the order of responses (main clause is begun first, followed by the parenthetical) doesn't match the order of questions responded to, a matter of discourse coherence rather than something attributable to the intrinsically NAI status of parentheticals.

These examples lead to two related conclusions: 1) we cannot hardwire NAI status into parentheticals—their variable (N)AI behavior must be derived; and 2) an account of parentheticals should explain why sentence-medial parentheticals display behavior characteristic of NAI content, while sentence-final parentheticals do not. This positional difference is acknowledged in Koev (2022), and we aim to explain it in our account as well.

2.2.2 Flexibly NAI accounts

The account in Koev (2013, 2022) is the only one we're aware of that avoids hard-wiring in NAI status to parentheticals. Unfortunately, it does so at the cost of carving up the empirical landscape—specifically, parentheticals like those in (1) and (4)—into 'pure' parentheticals that are easily accounted for by the theory, and 'impure' parentheticals, which require a distinct account. It also involves multiple null operators, which makes the theory not well-positioned to explain the cross-linguistically universal invocation of comma intonation to mark parentheticality.

For Koev, parentheticals are formed not with comma intonation, but with covert embedded illocutionary force heads ('P-FORCE') that separate them from the main assertoric content of the sentence. As far as we know, there is no overt representation of P-FORCE in any language (comma intonation presumably does not qualify, as it is represented twice for sentence-medial parentheticals, and also occurs in non-parenthetical expressions, as discussed above).

Because P-FORCE is an illocutionary force operator, it can by definition only combine with clauses, which shrinks the descriptive coverage of Koev's account: he can only use P-FORCE to explain the behavior of clausal parentheticals, not nominal ones like (4b) or (9), or adverbial ones, like (4a). These constructions receive distinct treatments in Koev's theory.

In the context of these 'pure,' clausal parentheticals, Koev weakens the traditional NAI categorization: he defines P-FORCE as marking parenthetical content as exempt from the at-issueness (e.g. QUD-addressing) requirements that ordinary context updates are subject to (making them "impoverished versions of root force operators", p129). This is shown in (12) for the parenthetical

declarative marker (there is also an interrogative version), for propositional content p and input context c .

$$(12) \quad \llbracket \text{P-FORCE}_c^p \rrbracket = \lambda P \exists p \wedge P(p) \wedge c + p$$

The conjunct ‘ $c + p$ ’ represents the automatic CG update; he specifies that because this one is formulated using $+$ and not $\oplus$, which is used for matrix assertions, there is no requirement that p be relevant to the CG; there is thereby (N)AI optionality built into the system.

Koev (2022) provides a distinct treatment for ‘impure’ parentheticals, i.e. any parenthetical whose content doesn’t denote a proposition, like the discourse adverbs in (4a). His account of discourse adverbs hard-wires the illocutionary status into the adverb (which incorrectly predicts these adverbs are unambiguously illocutionary; see Rett, 2021, for arguments against that assumption).

So Koev’s account doesn’t meet the target we’ve set for our own: his use of covert force heads i) falls short of standards for compositionality in relying on otherwise-unmotivated covert elements; ii) forces a disjunctive treatment of the class of parenthetical expressions; and iii) leaves unexplained why the signature feature of parentheticals is pauses, which do not plausibly track the distribution of covert force heads.

In what follows, we develop an analysis of parentheticality that contains no elements whatsoever. We argue that a fully general account of the meaning of pauses can capture the entire lay of the land: it can explain why pauses are necessary for parentheticality without predicting that all pauses should induce parentheticality; it can handle parentheticals of any size and of any relationship to the main clause they accompany; and it can explain why parentheticals display NAI behaviors in some contexts but not others.

3 A compositional account of comma intonation

Our proposal takes advantage of recent models that enter a representation of the speech event into the discourse record, and thus make it available for anaphora, drawing on the observation that lots of natural language anaphors can reference things in the world not explicitly linguistically introduced into the conversation. We suggest that comma intonation functions as a speech event indefinite, introducing a new speech-event discourse referent that may or may not be coreferent with others in the discourse. In this way, we can give comma intonation a single semantic analysis across all of its uses, and can also derive its optional not-at-issue status without relying on any covert operators.

3.1 Speech events and the start-up update

In a famous early discussion of discourse anaphora, Stalnaker observed:

“If a goat walked into the room, it would normally be presupposed, from that point, that there was a goat in the room. And the fact that this was presupposed might be exploited in the conversation, as when someone asks, ‘How did that thing get in here?’, assuming that others will know what he is talking about. In the same way, when I speak, I presuppose that others know I am speaking, even if I do not assume that anyone knew I was going to speak before I did. This fact, too, can be exploited in the conversation, as when Daniels says ‘I am bald’, taking it for granted that his audience can figure out who is being said to be bald.” (Stalnaker, 1978, 86)

This suggestion amounts to a treatment of the speech event, time, and speaker as ‘goat anaphora’: we can refer to these things not because they are explicitly introduced by the linguistic string (or by some covert operator or structure), but because they are self-evident in the discourse, like a goat that’s just stepped into the room. Since that proposal, dynamic theories of discourse anaphora and cross-sentential anaphora have adapted to be able to model discourse referents (drefs) that are introduced not just linguistically, but situationally (see also Egan and Sundell, 2022).

Bittner (2007) was the first to formalize this proposal of Stalnaker’s, and to use that formalism to compositionally model anaphora to entities like speaker and speech time. She proposed that the utterance of any sentence is first modeled with something called the **start-up update**, which introduces and focuses attention on three default topics: (i) the speech world w_0 ; (ii) the speech event e_0 (with its dependent variable, the speaker, or agent of the speech event); and (iii) the speech time t_0 . In her adaptation, Rett (2024) simplifies this, in the framework of CDRT (Muskens, 1995), as follows:

$$(13) \quad \text{the start-up update: } [e_0, x_0 \mid \text{speech}(e_0) \wedge \text{agent}(x_0, e_0)]$$

Informally, (13) can be read as an instruction to the discourse interlocutors: associate the present event of speaking with the variable e_0 , and its speaker with the variable x_0 . Crucially, this update is not intended to represent the semantic content of some covert operator, or some representation of an extra-sentential phenomenon like a speech act. It simply models, in a dynamic framework, the way that the onset of someone speaking makes salient that they’ve started speaking.

For Bittner (2007), the start-up update is used to formally and compositionally model first-person anaphora and present tense in polysynthetic languages like Kalaallisut. For Rett (2024), the start-up update is used to formally model illocutionary content (qua content about the utterance) as speech-event modification. She derives the difference in the interpretation of *briefly* between (14a) and (14b) in terms of proximity to the start-up update:

- (14) a. Briefly, she was the president.
- b. She was the president briefly.

In Rett’s approach, *briefly* modifies the most salient event (specifying that it is brief). In (14a), the most salient event for the sentence-initial *briefly* is the speech event itself, modeled via the start-up update’s introduction of a dref for that event. The result is an illocutionary interpretation of the adverb in which it characterizes the speech event itself. In (14b), in contrast, the most salient event is the topic event, introduced by the VP *was the president*. The result is a descriptive interpretation of that same modifier, in which it restricts the interpretation of the utterance. This use of the start-up update explains why *briefly* can be polysemous between an illocutionary and a descriptive meaning, and also why its interpretation is affected by its position in the sentence.

3.2 The core analysis: speech event indefiniteness

We borrow the start-up update from Rett (2024) (although we will ignore the role of the speaker dref, as it will generally be redundant in the modeling of parentheticals and other cases of self-interruption). But we depart from Bittner and Rett in specifying a larger distribution for the start-up update: rather than only applying at the beginning of an utterance (however long the utterance might be), the start-up update, for us, is initiated following any pause. Specifically, the start-up update kicks in as soon as someone begins talking after a break in speech, however brief. This is a natural extension of the proposal, as the start-up update models the fact that somebody

has started talking, and a pause is just somebody starting talking anew.³

Since comma intonation counts as such a break, we can specify the start-up update as a semantic contribution of the comma:⁴

$$(15) \quad \llbracket , \rrbracket = [e_0, x_0 \mid \text{speech}(e_0) \wedge \text{agent}(x_0, e_0)]$$

Before we present an example of (15) in action, we’ll introduce a few foundational, non-novel assumptions we make in the account.

First, we assume that proper names and indefinites introduce drefs, while pronouns and definites are anaphoric to them (Heim, 1982). We notate this as in (42) for the individual dref s .

$$(16) \quad \text{Sam}^s \text{ did his}_s \text{ exercises.}$$

Second, we assume that the relationship between a predicate in an appositive and its anchor (the entity it modifies) is discourse-anaphoric. That is, in the sentence in (17), the relative pronoun *who* receives its referent not because it’s in any syntactic relationship with the coreferring noun *Stephanie*, but because the dref s is the most salient entity available for *who* to refer to. This assumption is standard in a wide range of dynamic-semantic accounts of parentheticality (Sells, 1985; del Gobbo, 2007; Nouwen, 2007; AnderBois et al., 2013; Murray, 2014).

$$(17) \quad \text{Stephanie}^s, \text{ who}_s \text{ brought cookies, left after midnight.}$$

AnderBois et al. 2013 emphasize that this is a linear, but parallel, update, as we predict; (18) is theirs.

$$(18) \quad \text{I got a few quick words with [Halle Barry]}^h, \text{ who}_h \text{ wore Prada, and [Sigorney Weaver]}^s, \text{ who}_s \text{ didn't.}$$

Finally we assume that, while indefinites always introduce a new dref, these drefs can be coreferential. This is illustrated in (19), which also provides an illustration of how we’ll annotate the phenomenon of coreferential anaphora.

$$(19) \quad [\text{A linguist}]^x \text{ likes cake and } [\text{a linguist}]^{y(=x)} \text{ likes pie. (It's Stephanie!)}$$

In (19), the first indefinite introduces the dref x , and the second indefinite introduces the dref y . But, after proceeding context makes clear the two are coreferent, the speaker learns that $y = x$ (and, in fact, that $x = y = s$, although we don’t represent that here). For us, comma intonation—or pauses—are indefinite-like, but instead of introducing drefs corresponding to individuals, they introduce speech-event drefs.

³It should be easy to see that the same concept extends to cross-speaker interruptions. If D is talking, and J interrupts him, J’s utterance will carry its own start-up update, and introduce a new speech-event dref, despite there being no pause between D’s and J’s speech events. This is because J’s interruption comprises J beginning to speak, regardless of whether it was preceded by a pause in D’s utterance.

⁴Note again that we are not in the business of proposing semantic analyses of orthography—indeed, we don’t believe that any morpheme at all is involved in parentheticality. By putting a comma between denotation brackets, we’re using shorthand to represent the way that pausing and then continuing to speak introduces a new dref for the speech event following the pause.

3.3 Indefiniteness and speech-event anaphora resolution

3.3.1 The core account

The power of this account is demonstrated in the simple appositive in (20); we illustrate the relevant anaphora using the notation above, and provide a more concrete formal semantics below.

(20) Stephanie, who brought cookies, left after midnight.

- a. $(\text{pause})^{e_i} [\text{Stephanie}^s]_{e_i}$
- b. $,^{e_j} [\text{who}_s \text{ brought cookies}]_{e_j}$
- c. $,^{e_k} [\text{left after midnight}]_{e_k (=e_i)}$

(20) illustrates the core of our proposal: comma intonation—or, technically, the onset of speech following a pause—introduces a new speech-event dref, potentially corresponding to a distinct speech event. Hidden in the notation in (20) is our assumption, in the spirit of Bittner (2007) and Rett (2024), that semantic content is indexed in some way with the speech event it was introduced in, so that the conversational record can recover something like in (21):

- (21) e_i is a speech event with the content ‘s left after midnight’
 e_j is a speech event with the content ‘s brought cookies’

If a speech event is discontinuous, as in the case of a medial parenthetical, we assume that identifying the two speech events as coreferential entails merging the content of those speech events. There is a direct parallel to a sentence like *The mother drinks wine, and the professor drinks whiskey* in a context in which both definites are determined to be coreferential; at some level of semantic representation, the individual referred to must be represented as both a mother and a professor. For the sake of perspicuity, we are not formally representing this association between a speech event and its content, but we assume that interlocutors are.

Recall that pauses are optional in speech ((22) is repeated from (8)):

- (22) a. Nobody is both a good linguist and a good philosopher.
b. Nobody is both a good linguist, and a good philosopher.
c. Nobody, is both, a good linguist, and, a good philosopher.

3.3.2 Rejecting the pause-parenthetical biconditional

Given that all parentheticals are formed with pauses, but not all pauses form parentheticals, there are three theoretical options for addressing the role of pauses in the formation of parentheticals: i) we can ignore the role of pauses in parenthetical formation because there’s not a one-to-one correlation (Koev (2022) chooses this option); ii) we can argue that there is a difference between pauses that form parentheticals and those that don’t (this option is chosen explicitly by Truckenbrodt (2015) and implicitly by AnderBois et al. (2010, 2013), who only discuss pauses as they relate to the formation of parentheticals).

Or there is a third option: iii) analyze all pauses the same across constructions, but give them a relatively weak semantics, allowing them to contribute to the formation of a parenthetical or not, depending on context. We endorse this third option, largely because we believe that the conditional correlation (parenthetical $\rightarrow$ pause) is worth accounting for even in the absence of a biconditional (parenthetical $\leftrightarrow$ pause), especially given the apparently universal tendency of languages to co-opt pauses to mark parentheticals.

For us, this third option (and the weak “pause semantics” it requires) comes in the form of

the assumption that, like any indefinite, pauses are compatible with—but do not require—the new speech-event dref to be referentially distinct from a prior one. This marks a subtle but important change from the account in AnderBois et al. (2010, 2013), where the left comma toggles to not-at-issue content, and the right toggles back to at-issue content. We make this switch for several reasons: first, as previously discussed, parentheticals do not always encode not-at-issue content. Second, because of the optionality of pauses, discussed immediately above.

A final reason we analyze pauses as speech-event indefinites is because the use of parentheticals can result in more than two speech events, as illustrated in (23) (this is also observed in AnderBois et al., 2013).

- (23) $(pause)^{e_i}$ [Morris Halle^m] _{e_i}
 $,^{e_j}$ [who_m was a Leninist] _{p, e_j}
 $,^{e_k}$ [which_p I learned from reading his_m MA thesis] _{e_k}
 $,^{e_l}$ [was a respected linguist] _{$e_l (= e_i)$} .

The thing to note about (23) is that it involves several pairs of comma intonation, and at least some of them can’t be characterized as toggling from AI to NAI content: the comma superscripted with e^k separates one parenthetical from another. And while the second seems clearly subordinated to the first—the content encoded in e^k is modifying the content encoded in e^j —both parentheticals are NAI content in an out-of-the-blue context.

3.3.3 The role of prosody

We assume that content, context and prosody work together to resolve anaphora, i.e. to tell us whether one speech event is a continuation of a previous one. Arguably, the reason we interpret the final speech event in (23), e_l , as coreferential with the initial one (e_i), is attributable to several factors simultaneously: the content of each speech event (in particular, that neither forms a complete proposition, but that the latter can function as the completion of the former); the context of utterance (it's plausible that Morris Halle was a respected linguist); and the prosody with which all of the various speech events in (23) are pronounced.

Truckenbrodt (2015) (and many others) assumes there is a one-to-one correlation between pauses and IP boundaries; hence his generalization that IPs delineate speech acts. But IP boundaries are marked by *terminal contours* (Jun, 2022), and it seems like we use this intonation to help us determine whether a new speech event dref is distinct from, or coreferent with, a previous one.⁵ We observe that a pause after a completed terminal contour is compatible with that speech event being completed, while any pause prior to completion of a terminal contour is interpreted as having the potential for continuation.

- (24) a. H* H* L-L%
Dylan just finished eating, according to Tim.
b. H* H* L-L%
Dylan just finished eating, according to Tim, a muffin.

Concretely: (24b) illustrates that, despite containing two sentence-medial pauses, the sentence does not contain the three intonational phrases that one might expect. Rather, the main clause

⁵Sturman (2022) observes in both a corpus study and comprehension task that subclausal quotation, too, is obligatorily marked with an IP boundary at the right edge of the quoted material. This offers additional evidence that IP boundaries are used to demarcate speech events, and that that is something we know when we learn how to parse a language.

appears to be parsed as a single intonational phrase, containing a single terminal contour, despite being interrupted by the parenthetical material (which is parsed into a separate intonational phrase). In short, one prominent cue that tells us whether a thought is finished or interrupted is whether it carries a terminal contour. Clauses that are infixed with a sentence-medial parenthetical do not carry a terminal contour at either parenthetical boundary, and this reliably distinguishes them from uninterrupted clauses.

3.3.4 Indefiniteness and lexical competition

We end this subsection with a note about indefiniteness and lexical competition. While it's true that indefinites can be coreferential (as in (19)), they generally aren't; it's typically weird if not unacceptable to use coreferential indefinites, as demonstrated in (25).

- (25) ??[A woman]_x walked into a bar. [A woman]_{y(=x)} said hello. [A woman]_{z(=x)} ordered a beer.

But this is because indefinites compete lexically with definites. And once a discourse referent has been established, the definites are preferred as presuppositionally stronger, as demonstrated in (26) (see e.g. Heim's 1991 influential treatment of the phenomenon via the maxim 'Maximize Presupposition').

- (26) [A woman]_x walked into a bar. [The woman]_x said hello. She_x ordered a beer.

If indefinites weren't competing with definites, we'd expect (25) to be a legitimate way of expressing (26). And that's what we believe is happening with comma intonation: comma intonation is licit, and can have completely free coreferential variability, because there is no competing method of specifying a definite speech-event anaphorically.

To summarize: in our account, the start of speech after a pause introduces into the discourse record a new speech event discourse referent, one of several drefs we already know interlocutors keep track of when modeling the discourse for the purpose of anaphora and illocutionary content. Just as with other occasions of anaphora resolution, interlocutors use content, context and prosody to help determine which post-pause utterances are the continuation of prior ones (which amounts to a case of self-interruption), and which are not (which amounts to a parenthetical).

This perspective allows us to account for what appears to be a universal and obligatory strategy for marking parentheticals—a pause on either end—without committing to the false prediction that every pause necessarily marks the start or end of a parenthetical. While we model this semantic contribution of the pause compositionally in (15), this is intended as a representation of the introduction of a speech event as 'goat anaphora', instead of some sort of arbitrary linguistic convention: (15) simply models, in a dynamic framework, how the observation that the speaker has started talking again after a pause introduces a new dref for a speech event into the conversational model.

In the remaining section we discuss the empirical ramifications of this core account, and show how it provides natural solutions to a large set of puzzles identified in previous literature on parentheticals.

4 Empirical ramifications

Koev (2022) presents six empirical puzzles presented by parentheticals, and uses them as a list of

theoretical desiderata. Since we have already addressed our top theoretical priorities (the universality of comma intonation, and the assimilation to self-interruption), we spend this section going through other characteristics of parentheticals, and showing how our account derives them.

4.1 Semantic independence and scopelessness

We'll start with Koev's **embeddability puzzle** (i.e. semantic independence): parenthetical content is not part of the main assertion. We also address his closely related **projection puzzle**, i.e. the observation that parenthetical content can never scope under operators introduced elsewhere in the utterance. We presented evidence of the embeddability puzzle in (1) and (2), repeated below.

- (27) a. Andi said that a Canadian, who was talking about Kant all night, was the only philosopher there.
b. Andi said that a Canadian who was talking about Kant all night was the only philosopher there.

Koev illustrates the projection puzzle with (28), where the Minnesota restriction is under the scope of the possibility modal in (28b) but not (28a).

- (28) a. It's possible that Edward, who is from Minnesota, enjoys cold winters.
b. It's possible that Edward is from Minnesota and enjoys cold winters.

Both of these properties (semantic independence, scopelessness) are canonical properties of NAI content. These puzzles have traditionally been addressed (with the exception of Koev) by hard-wiring in the NAI status of parentheticals, but this is not what we do. For us, both of these facts are predicted by the independence of the two speech events associated with each proposition. To repeat the observation discussed in terms of (21), we predict that the parenthetical example in (28a) is interpreted as involving two distinct speech events, each with their own semantic content:

- (29) e_i is a speech event with the content 'It's possible e enjoys cold winters'
 e_j is a speech event with the content ' e is from Minnesota'

In our account, there is no syntactic relationship between the parenthetical utterance and the matrix utterance, so we expect complete semantic independence, just as we would in the corresponding (30) or even (31), where the two speech events have two different speakers.

- (30) It's possible that Edward ^{e} enjoys cold winters. He _{e} is from Minnesota.
(31) A: It's possible that Edward ^{e} enjoys cold winters.
B: He _{e} is from Minnesota.

To summarize: while semantic independence and scopelessness are indeed properties of NAI content, they are not exclusively properties of NAI content. They are also properties of separate utterances of separate sentences. In the case of the embeddability and projection puzzles, we account for the behavior of parentheticals by analyzing them as distinct utterances from the main-clause utterance.

4.2 Temporal precedence and the relation to the QUD

Another puzzle noted by Koev (2022) is the **at-issueness puzzle**: the observation that parenthetical content "cannot naturally answer the question under discussion" (p22-3). This is exemplified by (32) (from Koev, 2022).

- (32) A: Why are you so upset?
 B: Jon Stewart, who is my favorite comedian, just appeared on Fox News.
 B':#Jon Stewart, who just appeared on Fox News, is my favorite comedian.

The idea is that the B' response is unacceptable in the context because the content that addresses A's question is parentheticalized.

A related observation is that the content encoded in parentheticals cannot be addressed with standard, direct response particles, as demonstrated in (33) (in contrast with expressions like *Hey, wait a minute!*, which address backgrounded content, as discovered in von Fintel 2004).

- (33) A: Edward, who is from Minnesota, enjoys cold winters.
 B: #No he's not! He's from Michigan.
 B': Wait a minute, I thought he was from Michigan!

Both of these properties (inability to address the QUD, inability to be targeted by direct response particles) are also canonically associated with NAI content. Crucially for us, the two or more speech events associated with a parenthetical utterance (e.g. (21)) are not subordinated with respect to each other, and neither is demoted to some not-at-issue level or dimension. They are differentiated only in their temporal precedence, or order of introduction: in the acceptable response in (32), the speech event that addresses the QUD is temporarily primary to the one that does not. We argue that this temporal precedence is what actually underlies the apparent NAI behavior of parentheticals.

There is a conversational expectation that the QUD will be answered as soon as possible (Roberts, 1996). This is illustrated in (34) for the equative construction, whose two arguments are semantically equivalent (with respect to the predicate), but which nevertheless demonstrates topic sensitivity (Rett, 2015).

- (34) A: How tall is Hannah?
 B: Hannah is exactly as tall as Arthur.
 B':#Arthur is exactly as tall as Hannah.

In the context of parentheticals, this pressure for the temporally prior content to address the QUD translates to a main-clause association with at-issue content (i.e., content that addresses the QUD). The requirement that the main, temporally primary clause address the QUD leaves underspecified whether the parenthetical also addresses the QUD. And in fact, there is evidence that parenthetical can address the QUD when main-clause content does as well (see also (9)):

- (35) A: Why are you so upset?
 B: Jon Stewart just appeared on Fox News, where he was really friendly with the panelists, and he's my favorite comedian.

Furthermore, there are a number of environments in which it's been observed that embedded content can address the QUD as naturally as matrix content can; speech reports are one such case, as illustrated in (36) (Simons, 2007).

- (36) A: Where is Andrew?
 B: Ned said he's in Florida.

(37) recreates this phenomenon for parentheticals. The B response is an instance of a sentence-final parenthetical addressing the QUD, and the B' response illustrates the same phenomenon for a sentence-medial parenthetical.

- (37) A: Where is Andrew?
 B: As Ned said, he's in Florida.
 B': As Ned said, he's in Florida, which is not what I expected.

The conclusion we can arrive at is as follows: in an utterance consisting of more than one speech event, there is information-structural pressure for the QUD to be addressed as soon as possible, which means that the typical parenthetical will not address the QUD, as parentheticals are by definition not sentence-initial. But, as we've come to expect (Tonhauser et al., 2013), there is contextual variation in at-issue status: sentence-medial and sentence-final parentheticals can systematically address the QUD when the matrix clause is informationally relevant.

In terms of the ability of parenthetical content to be directly denied (33), we see that that generalization, too, is not without exception. Parentheticals can be responded to, but only if the speaker is immediately interrupted after the parenthetical, as in (38).

- (38) A: Stephanie, who brought cookies,—
 B: No she didn't!

This is evidence that response particles, too—or at least the direct ones, as opposed to the indirect *Hey wait a minute* ones—are temporally sensitive. In a standard, sentence-medial parenthetical like (28a), the matrix speech event both starts and ends the sentence. If we think of response particles as addressing the most salient proposition (which is quite a standard assumption Krifka, 2013), we predict that they are more likely to address sentence-final parentheticals, or interrupted parentheticals, as in (38). This is effectively the 'Right Frontier Constraint' from Asher and Lascarides (2003).

4.3 Discourse anaphora and definiteness restrictions

In this subsection, we set forth a number of observed syntactic and semantic differences between appositives and restricted relative clauses (de Vries, 2002), exemplified in (39), and explain them in the context of the proposed account.⁶

- (39) a. I spoke to the lecturers that failed the diacritic test. *restrictive relative clause*
 b. I spoke to the lecturers, who failed the diacritic test. *appositive relative clause*

Only restrictive relative clauses can begin with the complementizer *that*; appositives must begin with relative pronouns, like *who*. As we spell out in §3.2, in appositive relatives, this proform discourse-anaphorically refers to the referent made salient by the immediately preceding nominal.

The assumption that appositive relatives are interpreted by establishing an anaphoric relation between the relative pronoun and the discourse referent made salient by the immediately preceding nominal solves three other puzzles associated with parentheticals.

First, a canonical property of appositives is their non-restrictive status, and a theory of appositives as distinct utterances accounts for this as well. Restrictive modification requires compositional interaction; a restrictive relative clause denotes a property of individuals, which can intersectively modify the noun, which denotes the same kind of property. An appositive cannot involve this modification because, as an independent utterance, it does not interact compositionally with the nominal it follows. Rather, as stated above, the relative pronoun must be anaphorically linked with the discourse referent introduced by the precedent nominal.

⁶These examples come from de Vries (2006), who lists a third type of relative clause, a 'maximalizing' relative, which is just a degree relative, and must receive a distinct semantic interpretation from the other two (Rett, 2026).

Second, it follows directly that antecedents for appositives must be either definites, specific indefinites, or bound, because these are all and only the expressions that make discourse referents available for discourse anaphora. (40) demonstrates the restriction on specificity; (41) shows the bound case.

(40) *I saw someone or other, who wore a red hat.

(41) Every chess set comes with a spare pawn, which you will find taped to the top of the box.

In contrast, because appositives semantically restrict their pivot (via intersection, or Predicate Modification), there is no requirement that they be referential. It is just as easy for someone to talk about a professor who bowls as it is to talk about a professor, because we can concatenate indefinite modifiers in the old fashioned way. As a result, restrictive relatives are not subject to a definiteness restriction, but appositive relatives are.

Third, we can explain why appositives are generally restricted to post-nominal positions, as illustrated in (42).

- (42) a. Annie, who plays the violin, has bought a new bow. (cf. de Vries 2006 ex.3)
b. ?Annie has, who plays the violin, bought a new bow.

The oddness of (42b) can be explained by the appositive not being immediately adjacent to the nominal that makes salient the referent of its relative pronoun. It would be a poor utterance-planning choice on the part of the speaker not to utter the appositive at the point that referent is maximally salient, i.e., immediately following the nominal, and so (42b) can only be understood as the speaker catching themselves and making the late decision to add an appositive, and not as a fluent, planned utterance. Note that (42b) is improved in a context in which it's clear that it's an instance of self-interruption, i.e. utterances that aren't associated with efficient utterance planning:

(43) Annie has, who you met at my birthday party! I forgot!, been arrested.

Finally, our proposal can also explain why parentheticals can't be bound into by the main clause (Sells, 1985, Koev's **binding puzzle**). For us, parentheticals are completely distinct from the main clause, so we expect that binding into a parenthetical is just as unnatural as binding into a distinct sentence.

- (44) a. [Every student]^x greets the professor who remembers her_x.
b. *[Every student]^x greets Farkas, who remembers her_x.

Interestingly, because we treat the association of a parenthetical and its pivot as discourse-anaphoric, we predict that binding into parentheticals is subject to exactly the same restrictions as cross-sentential anaphora. And this is indeed what we see:

- (45) a. Every farmer owns [a cart]^x. He uses it_x for harvesting.
b. Every farmer owns [a cart]^x, which_x he uses for harvesting.
(46) a. No farmer owns [a cart]^x. *He uses it_x for harvesting.
b. No farmer owns [a cart]^x, *which_x he uses for harvesting.

4.4 Self-contradiction and the parenthetical Moore's Paradox

Our proposed solution to the at-issueness puzzle, especially as it pertains to the licensing of response particles, is clearly related to Koev's **hedging puzzle**: the observation that "modification by a

slifting parenthetical reduces the possibility of hedging as compared with modification by a matrix clause” (p23). This is illustrated in (47) (from Asher, 2000).

- (47) a. Mary assures us that John can be trusted, but I don’t trust him.
b. #John, Mary assures us, can be trusted, but I don’t trust him.

In our account, each of these pauses contributes a potentially distinct speech-event dref, and content, context and prosody conspire to piece together speech-event coreference in a particular way, namely as shown in (48).

- (48) a. (for (47a):)
 e_i is a speech event with the content ‘Mary^m assures us that John^j can be trusted’
 e_j is a speech event with the content ‘But I don’t trust him_j’
b. (for (47b):)
 e_i is a speech event with the content $p = \text{‘John}^j \text{ can be trusted’}$ ⁷
 e_j is a speech event with the content ‘Mary^m assures us that p ’
 e_k is a speech event with the content ‘But I don’t trust him_j’

Given that all three speech events in (48b) have the same speaker, and given that speaker belief that p is apparently a precondition on the sincere assertion of p , the content of e_i and the content of e_k are contradictory. And we know independently that it’s infelicitous to explicitly state disbelief in an assertion you’ve previously committed to; this is Moore’s Paradox (Moore, 1993), standardly exemplified in (49).

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

Koev was careful to phrase the hedging puzzle in terms of a tendency rather than an absolute (slifting “reduces the possibility of hedging”, *ibid.*), and for good reason: when slifting doesn’t result in self-contradiction, hedging is totally acceptable, as demonstrated in (50).

- (50) John, Mary assures us, can be trusted, but I don’t trust Mary.

Finally, Koev also discusses the **polarity puzzle**, which he characterizes as a prohibition against slifting with upward-monotonic embedders (Jackendoff, 1972).

- (51) a. John is, I {think / *don’t think / *doubt }, a fink.
b. John is, I don’t doubt, a fink.

For us, (51) consists of two contentfully distinct speech events, listed in (52) (where e_j represents the speech events associated with a positive epistemic relation to p , and $e_{j'}$ represents the speech events associated with a negative epistemic relation to p).

- (52) e_i is a speech event with the content $p = \text{‘John is a fink’}$
 e_j is a speech event with the content ‘I {think / don’t doubt} p ’
 $e_{j'}$ is a speech event with the content ‘I {don’t think / doubt} p ’

As before, the speaker of e_i is identical to the speaker of e_j , and the result is a contradiction between

⁷Note that we will need to invoke propositional anaphora (Krifka, 2013; Snider, 2017) in a formal account of this utterance: namely that a proposition p is added to the top of the Stack, which models salience, as soon as it is uttered in completion. Just like other cases of discourse anaphora, we expect salience and anaphoric availability to be linearly governed, which means that the anaphoric invocation of p in e_j is related to the p introduced by e_i discourse-anaphorically, as explained in §3.2.

the felicity conditions on the assertion of the content of e_i and the content of e_j .

So, from our perspective, Koev’s hedging and polarity puzzles are just parenthetical versions of Moore’s Paradox: since there is no difference in the type of content introduced by e_i and e_j , just a temporal precedence difference, we can assimilate the content introduced in e_j to an assertion, subjected to whatever governs Moore’s Paradoxicality.

Our takeaway for this subsection is that there is a species of puzzle that appears to be rooted in the idea that a slifting parenthetical ultimately amounts to something semantically equivalent to a single attitude report—*It’s raining, Mary assures us* means the same thing as *Mary assures us that it’s raining*—rendering differences in behavior between slifting parentheticals and the corresponding attitude reports mysterious. But for us, slifting parentheticals are equivalent to an attitude report combined with an assertion of its prejacent—*It’s raining, Mary assures us* means the same thing as asserting both *It’s raining* and *Mary assures us it’s raining*—which renders the purportedly puzzling differences between shifts and unshifted attitude reports unmysterious.

4.5 Other parenthetical issues

In this section, we discuss several new empirical predictions this account makes that, to our knowledge, no extant compositional account of parentheticality can make (at least straightforwardly).

In contrast to other accounts, we correctly predict the existence of stacked parentheticals, i.e. parentheticals that introduce more than two propositions. This was exemplified in (23) (reproduced in (53)).

- (53) [Morris Halle]^m, [who_m was a Leninist]^p, which_p I learned from reading his_m MA thesis, was a respected linguist.

While many don’t remark on the existence of these stacked parentheticals, it is explicitly ruled out by the account in AnderBois et al. (2010, 2013), which defines only left and right commas, which relate only parenthetical content to non-parenthetical content.

And, in contrast to Koev (2022), we can account for the fact that parentheticals need not have assertoric or interrogative illocutionary mood (or even clausal content).

- (54) a. That Canadian philosopher—get off my foot!—is still talking about Kant.
b. That Canadian philosopher—ouch, my foot!—is still talking about Kant.

We can explain why sequences of a main clause and a parenthetical interact with QUDs in a manner parallel to sequences of conjuncts in a conjunction (Kroll and Rysling, 2019; Syrett and Koev, 2015) because, for us, the difference between the matrix speech event e_i and the parenthetical speech event e_j is just one of temporal precedence. And we can explain why parentheticals appear to be processed independently of the sentences containing them (Dillon et al., 2014) because, in our proposal, they are distinct speech events, as semantically and illocutionarily independent as two successive assertions.

Finally, we are in a position to explain why sentence-medial speech modifiers can only modify the present speech event when they’re parentheticalized, as observed in Rett (2024).

- (55) a. Briefly, she was once president. *illocutionary only*
b. She was once briefly president. *descriptive only*
c. She was once, briefly, president. *ambiguous*

It’s challenging to account for a single modifier having both an illocutionary reading (in which it modifies the present speech event) and a descriptive one (in which it modifies the topic event),

but it’s especially hard to account for the availability of both readings in a single utterance. With the right model of salience, an account of this phenomenon is well within reach: we only need to be able to model the fact that, in the middle of the sentence, both the speech event and the topic event are equally salient. To do this, we’d need to rely on the idea (like that explicitly proposed in Krifka, 2014) that a topic event is introduced by the tense marker (in the TP), and only overrides the salience of the speech event at or near the end of the utterance. We leave such attentional issues to a cognitive model of salience (as argued for in Stojnić et al., 2017).

5 Conclusions & extensions

When you pause, either before a sentence or within it, you signal to your interlocutors that you might be starting a new speech event, and there is independent reason to believe language users keep track of different speech events (Bittner, 2007; Rett, 2024). Your interlocutors can use several cues—context, content, and prosody—to help determine whether a new speech event is distinct from a previous one, or whether you’re continuing to form a prior one. They do this whether the interruption seems relevant (traditionally called a parenthetical) or whether it doesn’t (what we’ve called self-interruption). We model this compositionally and discourse-anaphorically by treating all pauses as introducing new speech-event drefs. In this way, they operate as illocutionary indefinites.

We derive Truckenbrodt’s (2015) previously stipulative distinction between optional and obligatory pauses. In our account, a single speech event can be done all in one go or, equivalently, split across multiple pauses that all refer back to it. But two distinct speech events require two distinct start-up updates, necessitating a pause between them; so pauses are obligatory for parentheticals, but do not obligatorily induce parentheticality. We believe this is the most advantageous component of the account: we assign to pauses (or comma intonation) a meaningful role, but correctly predict a correlation that is crucially only unidirectional: if it’s a parenthetical, it contains a pause, but if it contains a pause, it’s not necessarily a parenthetical.

Related, because we present a natural, independently-motivated treatment of what a pause does semantically (or, more accurately, what speech immediately following a pause does), we correctly predict that parentheticality is marked the same way across languages. Our account relies on no null operators, so we don’t have to worry about acquisition issues, or lexical variation in parentheticality marking. We have no construction-specific operators, so we don’t have to worry about defending subclasses of parentheticals, or even of cases of self-interruption. And since pauses are a natural part of language—spoken or signed—we predict the absence of (obligatory) parenthetical morphology cross-linguistically.

Also related, we naturally explain the phenomenon of self-interruption, which is marked identically to parentheticals cross-linguistically.

(56) Stephanie, oh my gosh I’m so hungry!, brought some cookies.

Our proposal avoids hard-coding the NAI status of parentheticals into an analysis of their discourse contribution; a parenthetical, on our account, is an utterance like any other, distinguished only by the fact that it runs in parallel to another utterance, instead of being linearly sequenced. We suggest that the variable NAI behavior of parentheticals can be derived entirely from discourse considerations on the temporal ordering of utterances, with cases in which parentheticals are not available for polarity particle anaphora or cannot address the QUD reducing to situations in which, due to temporal precedence relations, they are not the most salient target for propositional anaphora, or do not comprise orderly ways to address the current question.

This reliance on temporal precedence to explain the acceptability of prosodically separate ma-

terial may also prove fruitful as an approach to analyzing the combination of prosody, discourse prominence, and temporal displacement at issue in nearby phenomena like topicalization:

(57) That book, Mary said she was going to read.

We leave a fuller investigation of such parallels to future work.

References

- AnderBois, S., Brasoveanu, A., and Henderson, R. (2010). Crossing the appositive/at-issue meaning boundary. In Li, N. and Lutz, D., editors, *Proceedings of SALT 20*, pages 328–346. CLC. doi: 10.3765/salt.v20i0.2551.
- AnderBois, S., Brasoveanu, A., and Henderson, R. (2013). At-issue proposals and appositive impositions in discourse. *Journal of Semantics*. doi: 10.1093/jos/fft014.
- Asher, N. (2000). Truth conditional discourse semantics for parentheticals. *Journal of Semantics*, 17:31–50. doi: 10.1093/jos/17.1.31.
- Asher, N. and Lascarides, A. (2003). *Logics of conversation*. Cambridge University. doi:.
- Bittner, M. (2007). Online update: temporal, modal, and *de se* anaphora in a polysynthetic discourse. In Barker, C. and Jacobson, P., editors, *Direct compositionality*, pages 363–404. Oxford University Press. doi: 10.1093/oso/9780199204373.003.0011.
- Dillon, B., Clifton Jr., C., and Frazier, L. (2014). Pushed aside: parentheticals, memory, and processing. *Language and Cognitive Neuroscience*, 29:483–98. doi: 10.1080/01690965.2013.866684.
- Egan, A. and Sundell, T. (2022). All update is goat update. Ms., Rutgers University and the University of Kentucky.
- Esipova, M. (2018). Constraints on QUD-addressing appositives: Pragmatics or prosody? Ms, NYU.
- Extepare, R. (1997). *The grammatical representation of speech events*. PhD Thesis, University of Maryland.
- von Fintel, K. (2004). Would you believe it? the King of France is back! In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 315–341. Oxford University Press. doi: 10.1093/oso/9780199270514.003.009.
- del Gobbo, F. (2007). On the syntax and semantics of appositive relative clauses. In Dehé, N. and Kavalova, Y., editors, *Parentheticals*. John Benjamins.
- Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. PhD Thesis, University of Massachusetts, Amherst.
- Heim, I. (1991). Artikel und definitheit. In von Stechow, A. and Wunderlich, D., editors, *Semantik: ein internationales Handbuch der zeitgenössischen Forschung*, pages 487–535. de Gruyter.
- Jackendoff, R. (1972). *Semantic interpretation in generative grammar*. MIT Press.
- Jun, S. (2022). The ToBI transcription system: Conventions, strengths, and challenges. In Barnes, J. and Shattuk-Hufnagel, S., editors, *Prosodic Theory and Practice*, pages 151–181. MIT Press.
- Koev, T. (2012). On the information status of appositive relative clauses. *Lecture notes in computer science*, 7218:401–410.
- Koev, T. (2013). *Apposition and the structure of discourse*. PhD Thesis, Rutgers University.
- Koev, T. (2022). *Parenthetical meaning*. Oxford University Press.
- Krifka, M. (2013). Response particles as propositional anaphora. *Semantics and Linguistic Theory*, 23:1–18.
- Krifka, M. (2014). Embedding illocutionary acts. In Roeper, T. and Speas, M., editors, *Recursion, complexity in cognition*, pages 125–155. Springer.
- Kroll, M. and Rysling, A. (2019). The search for truth: appositives weigh in. *Semantics and Linguistic Theory*, 29:180–200.
- Moore, G. (1993). Moore’s Paradox. In Baldwin, T., editor, *G.E. Moore: Selected Writings*, pages 207–212. Routledge.
- Murray, S. (2010). *Evidentiality and the structure of speech acts*. PhD Thesis, Rutgers University.
- Murray, S. (2014). Varieties of update. *Semantics & Pragmatics*, 7:1–53.
- Muskens, R. (1995). Tense and the logic of change. In Egli, U., Pause, P., Schwarze, C., von Stechow, A., and Wienold, G., editors, *Lexical knowledge in the organization of language*, pages 147–184. John Benjamins.

- Nouwen, R. (2007). On appositives and dynamic binding. *Research on Language and Computation*, 5:87–102.
- Potts, C. (2005). *The logic of conventional implicature*. Oxford University Press.
- Rett, J. (2015). *The semantics of evaluativity*. Oxford University Press.
- Rett, J. (2021). The semantics of emotive markers and other illocutionary content. *Journal of Semantics*, 38:305–340.
- Rett, J. (2024). Characterizing illocutionary content. *Semantics and Linguistic Theory (SALT)*, 34:685–712.
- Rett, J. (2026). The semantics of degree relatives. *Annual Review of Linguistics*, 12:61–80.
- Roberts, C. (1996). Informative structure in discourse: towards an integrated formal theory of pragmatics. In Yoon, J. and Kathol, A., editors, *OSU Working Papers in Linguistics*, volume 49. OSU.
- Searle, J. and Vanderveken, D. (1985). *Foundations of illocutionary logic*. Cambridge University Press.
- Selkirk, L. (1995). Sentence prosody: intonation, stress, and phrasing. In Goldsmith, J., editor, *The Handbook of Phonological Theory*, pages 550–569. Blackwell.
- Sells, P. (1985). Restrictive and non-restrictive modification. *CSLI*, 85-28.
- Sharf, E. (2024a). Restrictive modifiers in parenthetical positions. *NELS*, 54.
- Sharf, E. (2024b). What appositives can tell us about names and definite descriptions. In *Proceedings of Amsterdam Colloquium*.
- Simons, M. (2007). Observations on embedding verbs, evidentiality and presupposition. *Lingua*, 117:1034–1056.
- Snider, T. (2017). *Anaphoric reference to propositions*. PhD thesis, Cornell University.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9:315–332.
- Stojnić, U., Stone, M., and Lepore, E. (2017). Discourse and logical form: pronouns, attention and coherence. *Linguistics & Philosophy*, 40:519–547.
- Sturman, B. (2022). *Subclausal quotation intonation*. PhD Thesis, UCLA.
- Syrett, K. and Koev, T. (2015). Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics*, 32:525–577.
- Tonhauser, J., Beaver, D., Roberts, C., and Simons, M. (2013). Towards a taxonomy of projected content. *Language*, 89:66–109.
- Truckenbrodt, H. (2015). Intonation phrases and speech acts. In *Parenthesis and ellipsis*, pages 301–349. De Gruyter Mouton.
- de Vries, M. (2002). *The syntax of relativization*. LOT.
- de Vries, M. (2006). The syntax of appositive relativization: on specifying coordination, false free relatives, and promotion. *Linguistic Inquiry*, 37:229–270. doi: 10.1162/ling.2006.37.2.229.
- Wang, L., Reese, B., and McCready, E. (2005). The projection problem of nominal appositives. *Snippets*, 10:13–14.