

Composing Discourse Parenthetical Reports¹

Julie HUNTER — *Universitat Pompeu Fabra, Barcelona and Université Paul Sabatier, Toulouse*
Nicholas ASHER — *IRIT, CNRS, Toulouse*

Abstract. Hunter (2016) proposed that a speech report with a parenthetical interpretation but non-parenthetical syntax will contribute a modal discourse relation of the form $\Diamond R$ to discourse logical form. This paper provides a compositional account of the mechanism by which these modal relations are triggered. It then extends Hunter’s proposal to reports involving factive embedding verbs and provides an explanation of why explicit discourse connectives sometimes block parenthetical readings of reports in their scope.

Keywords: coercion, discourse connectives, discourse structure, factive verbs, parenthetical reports, speech reports

1. Introduction

As observed in Simons (2007) and Urmson (1952) *inter alia*, sometimes the main point of a speech or attitude report is conveyed by the embedded clause alone. Compare (1) and (2):

- (1) a. Liz is sad.
b. **A famous critic said that her new album isn’t worth a dime.**
- (2) a. John didn’t come to my party.
b. Jill said that **he was out of town.**

In both (1) and (2), (b) is naturally understood as offering an explanation of the content of (a); the intuitive explanantia are in bold. In (1), it is the fact that the critic said what she did that explains why Liz is upset; in other words, the entire report contributes to the explanation. On the most natural reading of (2), however, the embedded clause of (2b) alone contributes to the explanation; the fact that Jill said what she did plays a backgrounded, evidential role.

Urmson (1952) called uses of embedding verbs like that in (2) *parenthetical*. In these uses, embedding verbs play a secondary role, apparently serving to provide evidence for, or qualification of, the reported content. In this way, they resemble verbs of speech or attitude in the slifted clauses (Ross, 1973) of syntactically parenthetical reports such as (3).

- (3) He was out of town, Jill said.

Despite the similarity between (2b) and (3), we will assume that (2b) has a parenthetical reading (because *say* has a parenthetical reading) but non-parenthetical syntax. That is, following Simons (2007), we will assume that in (2b), the content expressed by the clause in bold falls in the scope of the embedding verb *say*, as the surface form of the report suggests.

¹We gratefully acknowledge research support from ERC research grant 269427.

Hunter (2016) developed a discourse-based account of reports with parenthetical readings but non-parenthetical syntax. In particular, Hunter proposed that the discursive contribution of a parenthetical report is to introduce a modal rhetorical relation between the embedded content of the report and the discourse preceding the report. In this paper, we build on the account of coercion developed in Asher (2011) and Asher and Luo (2012) to provide a compositional analysis of Hunter's proposal. We then extend the resulting account to model two facts about parenthetical reports not covered in other models: the fact that factive verbs such as *found out* can also be used parenthetically, as observed in Simons (2007), and the fact that certain discourse connectives seem to block parenthetical readings of reports in their syntactic scope, as discussed in Hunter and Danlos (2014). (4b), for instance, only has a non-parenthetical reading.

- (4) a. John didn't come to my party
b. because Jill said he was out of town.

We begin in Section 2 with a presentation of Hunter (2016)'s proposal. We then present the basics of Asher and Luo's account and our application of it in Section 3. Section 4 extends the resulting compositional account to parenthetical readings of factive verbs, and Section 5 examines the effects of explicit discourse connectives on the interpretation of reports.

2. The function of discourse parenthetical reports

Following Hunter (2016), we will call reports with a parenthetical interpretation but non-parenthetical syntax, such as (2b), *discourse parenthetical* to indicate that their parenthetical interpretation follows from discursive facts. There are no morphological features, like the evidential markers found in many languages, that force (2b) to receive an evidential reading. Nor should we assume that (2b) has an underlying parenthetical syntax that differs from its surface form. (See Simons, 2007, for arguments against a syntactic account.) There may be semantic differences between a parenthetical reading of a report and a non-parenthetical reading of the same report in the sense that the two interpretations might support different entailments. Normally, an instance of a speech or attitude report with a non-factive embedding verb such as *say* does not entail any level of commitment to the embedded clause, as illustrated by (1). In the absence of further context, however, the discourse in (2) arguably entails the possibility that John was out of town: if the speaker is not committed to this possibility, the relevance of (2b) to (2a) is unclear. Still, these entailments do not follow from (1b) and (2b) alone; they are *discourse-level* entailments that result from the interaction of a report and another discourse unit in the preceding discourse.

That the interpretation of a discourse parenthetical report results from the way a report is used in a discourse has been observed numerous times in the literature on parenthetical reports in English (Hooper, 1975; Simons, 2007; Urmson, 1952). Formal models of their discursive behavior are harder to come by. What does it mean to say that a parenthetical reading follows from the way that a report is used? How do we define the discourse function of a parenthetical report?

Simons (2007) sketched a diagnostic for parenthetical readings using question/answer pairs like (5).

- (5) a. Why wasn't John at the party?
b. Jill said that he was out of town.

In (5), only the embedded clause of (5b) is 'at-issue', in Simons' sense, because only the embedded clause provides an answer to the question posed in (5a). According to Simons' proposed diagnostic, a parenthetical report is one in which only the embedded clause is at-issue. Thus this diagnostic predicts, seemingly correctly in this case, that (5b) is parenthetical.

This suggestion is not yet an account, however, because it does not provide a general model of discourse function that could be applied in the absence of explicit question/answer pairs. In the case of (2), we can imagine that (2b) answers a question like *Why not?*, but short of a predictive account of how one can identify implicit questions, Simons' question/answer-based diagnostic cannot be used for arbitrary discourses. Moreover, the diagnostic does not yield an analysis of cases in which a speaker explicitly denies the embedded content of a seemingly parenthetical report. Suppose, for example, that the speaker in (2) followed (2b) with (6):

- (6) But that's a total lie because I saw him out walking his dog early the next morning.

Given (6), we cannot say that the embedded content of (2b) answers the question *Why not?* or *Why didn't John come to your party?*—the speaker explicitly commits herself to the negation of that content.

A final concern is how to model reports for which both the embedded clause and the report as a whole are relevant to the discourse but in different ways. To put this in terms of Simons' diagnostic, how do we handle cases in which both the embedded clause and the report as a whole answer questions that have been posed in the discourse, as in (7)?

- (7) a. Why didn't John come to your party? Did anyone tell you?
b. Jill said he was out of town.

In such examples, it's hard to say that one part of the report is at-issue while the other isn't; it's rather that the embedded clause and the report as a whole are *at-issue relative to different discourse units*. Of course, in (7), the second question is obviously very relevant to the first, but as emphasized in Hunter (2016), this is not always the case; the two parts of a report can be relevant to separate discourse units whose relevance to one another is much less clear.

Hunter (2016) argues that rhetorical theories (Asher, 1993; Asher and Lascarides, 2003; Hobbs, 1985; Mann and Thompson, 1988) can provide a model for the discourse function of parenthetical reports that generalizes Simons' suggestion and is more suitable for arbitrary discourses.²

²This suggestion has been made before. See Danlos and Rambow (2011); Dinesh et al. (2005); Hardt (2013); Hunter et al. (2006). Hunter (2016) examines past efforts to model parenthetical reports in rhetorical accounts and details many of the problems faced by these earlier proposals.

In a theory of rhetorical structure, what it means to have a certain discourse role or function is to enter into a particular rhetorical relation (or set of relations). The heart of Hunter (2016)'s proposal is that a report p is parenthetical just in case the embedded clause of p enters into a rhetorical relation R with a discourse unit that is discourse prior to p —that is, prior to the entire main clause of p —and the attributive predicate does not contribute to this relation. In (2), for example, (2a) is discourse prior to (2b), but the attributive predicate of (2b) does not contribute to the explanatory relation that is proposed to hold between the content in its scope and (2a):

(8) Explanation(*John didn't come to my party, he was out of town*)

This approach bypasses the need for implicit questions³, allows that both the embedded clause and the report as a whole can play independent rhetorical roles, and generalizes Simons' diagnostic for parenthetical reports on the assumption that answering a question yields a rhetorical relation between a question and its answer.

An important aspect of Hunter's proposal is that the rhetorical relation that relates the embedded clause of a parenthetical report to the preceding discourse will fall in the scope of a possibility operator. (8) poses a problem as it stands because Explanation is a *veridical* relation (Asher and Lascarides, 2003). This means that an instance of (8) cannot be true in a context unless both of its arguments are true in the same context. This is intuitively correct for (2): for John's being out of town to explain his absence at the party, he must have actually been out of town. The logical form in (8) is therefore too strong; as stressed by Simons (2007), reports are often used in situations like (2) precisely to hedge one's commitment to the embedded content. In some contexts it might be possible to infer full speaker commitment to this content, but a parenthetical use of a report does not in and of itself require such a strong commitment.

To weaken the entailments of a parenthetical report, the first step is to exploit the syntactic and semantic features of the report and assume that the attributive predicate takes not only syntactic scope over the embedded clause, but discourse scope over it as well. That is, the embedded clause is related to the attributive predicate via a rhetorical relation of Attribution, and the embedded clause is subordinate to the attribution predicate in the discourse representation. (9) gives the contribution of the Attribution relation to the logical form of the discourse.

(9) Attribution(*Jill said that, he was out of town*)

Semantically, this means that the embedded content of a report will not be entailed, at least not in virtue of figuring in an Attribution relation. (10) describes the entailments of the two arguments of an Attribution relation.

- (10) a. *Jill said that*: $\models \exists p.\text{said}(\text{Jill}, p)$
 b. *he was out of town*: specifies the content of p , but $\not\models p$

Of course, the Attribution relation is a general one that will apply to non-parenthetical reports

³We leave open here the possibility that discourse relations might be modelled in a question-based model, though see Hunter and Abrusán (2016) for a discussion.

as well. As such, Attribution will not support the stronger entailments that we see with parenthetical reports. Thus we arrive at the following impasse: if in building the representation for a parenthetical report, we ignore the discursive contribution of the Attribution relation, we will predict overly strong entailments. If we add Attribution to correct this, we end up with overly weak entailments.

To get around this problem, Hunter (2016) proposed a general rule that any time a discourse unit β figuring in an Attribution relation of the form $Attribution(\alpha, \beta)$ is rhetorically related to a unit γ that is textually prior to α , and thus outside of the scope of the Attribution, a relation R that can be inferred to hold between γ and β based on the contents of γ and β alone is weakened to $\Diamond R$. This rule is formalized in PA below, where e and e' are edges in a discourse graph that connect nodes representing discourse units; α, β and γ are labels for discourse units; and l is a labelling function that maps each edge to a label for a discourse relation (e.g., Attribution).

$$PA: (\exists e.e(\alpha, \beta) \wedge l(e) = Attribution) \rightarrow (\exists \gamma \exists e'(\gamma <_t \alpha \wedge e'(\gamma, \beta)) \rightarrow \exists R(l(e') = \Diamond R))$$

The effects of PA are illustrated in Figure 1.




Figure 1: Relations to units inside Attributions

PA is designed to capture the idea that when a speaker uses a report parenthetically, she uses the attributive predicate as a kind of buffer so that the embedded content can be relevant, but she does not have to fully commit to its content. From the perspective of discourse structure, PA reflects the fact that even if an attribution predicate does not seem relevant or ‘at-issue’ with regard to the preceding discourse, it nevertheless plays a crucial function in the overall discourse by affecting the relation inferred between its embedded clause and the preceding discourse.

Some consequences of adopting PA are, first, that it respects the constraint of veridicality imposed by many discourse relations, including Explanation, because a modal discourse relation only entails the possibility that its arguments are true:

$$(11) \quad \Diamond R(\alpha, \beta) \models \Diamond(\alpha \wedge \beta)$$

A full commitment to the content of β will be consistent with a formula of the form $\Diamond R(\alpha, \beta)$, but it will not be necessary. At the same time, commitment to $\neg\beta$ will be inconsistent with a formula of the form $\Diamond R(\alpha, \beta)$, as desired.

A further desirable consequence is that both α and β can be discourse relevant. Nothing in this proposal hinges on a binary distinction between at-issue and not-at-issue content, which, as argued in Hunter (2016), is untenable when we look at the way reports are actually used in longer stretches of discourse.

Hunter’s account, however, leaves some important questions about the nature of PA unanswered. First, what triggers PA and can we derive it compositionally? Also, does $\diamond R(\alpha, \beta)$ yield the right entailments for α ? Are there any examples in which a parenthetical report weakens a speaker’s commitment to the argument α ? In the next section, we argue that the transformation described by PA is a type of coercion, as understood in the frameworks of Asher (2011) and Asher and Luo (2012), and explain the entailments of a formula of the form $\diamond R(\alpha, \beta)$ in more detail.

3. Coercion

Coercion is an observed process whereby the meaning of a predicate P combines with the meaning of an expression e in its scope to produce a meaning shift. To illustrate, the predicate *is a very fancy and tasty Bordeaux* in (12) combines with the noun phrase *this bottle*, yet the final interpretation does not involve a predication of the bottle, but of the *contents* of the bottle.

(12) This bottle is a very fancy and tasty Bordeaux.

The combination of *is a very fancy and tasty Bordeaux* with *the bottle* coerces a shift in meaning. Coercion, many have argued, is the result of an adjustment to either the meaning of the predicate, the meaning of the argument, or to the composition that fits their meanings together, so that the argument of the predicate satisfies the predicate’s selectional restrictions, which in “ordinary circumstances” the argument does not. In (12) the predicate has to hold of wine or at least something that is edible or potable, and a mechanism of coercion like that proposed in Asher (2011) reinterprets the predication so that the predicted meaning is something like *the content of this bottle is a very fancy and tasty Bordeaux*.

Our claim is that a relation of the form $\diamond R$ that relates the embedded content of a parenthetical report to the preceding discourse is the product of coercion, triggered by a conflict between the demands of the Attribution relation and the demands of R . To unpack this claim, we begin with a brief overview of some relevant aspects of *Segmented Discourse Representation Theory* (SDRT; Asher and Lascarides, 2003), the formal framework that we will use to develop our coercion account.

The language of SDRT contains a countable set of discourse unit labels $DU = \{\pi, \pi_1, \pi_2, \dots\}$, and a finite set of discourse relation symbols $\mathbb{R} = \{R, R_1, \dots, R_n\}$, and formulas ϕ, ϕ_1, \dots from some fixed language L for describing elementary discourse move contents, where L is a language like that of higher order logic used in, for instance, Montague Grammar. SDRT formulas are of the form $\langle \pi : \phi \rangle$, where ϕ is either: (i) a formula of L ; (ii) a relational formula of the form $R(\pi_1, \pi_2)$, which says that π_1 stands in relation R , e.g., Explanation or Attribution, to π_2 ; (iii) a formula of the form $\diamond \psi$ where ψ is an SDRT formula; or (iv) a conjunction of SDRT

formulas. To provide an illustration, (2) yields three discourse units, π_1 , π_2 , and π_3 :

π_1 : John didn't come to my party.

π_2 : Jill said that

π_3 : he was out of town.

Each of these units is then associated with a formula of L — ϕ_1 , ϕ_2 and ϕ_3 , respectively—that specifies its content. For ϕ_2 , for example, the formula is $\exists p.\text{said}(\text{Jill}, p)$; for ϕ_3 , the formula is, $\text{Out-of-town}(\text{Sel}(\text{John}, \text{Jill}, \text{party}))$, where Sel , a part of the vocabulary of L , is a function picking out an appropriate individual type variable or discourse referent from the logical forms for the discourse so far (the formulas for π_1 and π_2). The combination of $\pi_1 - \pi_3$ in (2) will also yield two further formulas, $\text{Attribution}(\pi_2, \pi_3)$ and $\diamond \text{Explanation}(\pi_1, \pi_3)$.

How do we derive these formulas? While the Attribution relation between π_2 and π_3 is syntactically determined and triggered by the embedding verb *say*, the derivation of $\diamond \text{Explanation}(\pi_1, \pi_3)$ involves a more complicated process. SDRT requires that any discourse unit π' must, if it is not discourse initial, be rhetorically connected to some other unit π , if it is to make a coherent contribution to the discourse in which it figures; every coherent discourse can be represented as a directed, acyclic and weakly connected graph. This means that at least one part of the report in (2b) must be related to the preceding discourse. In this case, the preceding discourse consists of a single unit, so the task is simplified: either π_2 or π_3 must be related to π_1 .⁴

Often, the relation to the preceding discourse is not specified via a syntactic connection or an explicit discourse connective such as *because*. Where the selection of a relation or a discourse constituent is left unspecified, SDRT exploits the Sel function; $\text{sel}_{\mathbb{R}}(\text{sel}_{\Pi}\{\pi_m, \dots\}, \text{sel}_{\Pi}\{\pi_n, \dots\})$, for instance, says that some relation from the set \mathbb{R} of relations must be selected to apply to a selection of constituents from $\{\pi_m, \dots\}$ and from $\{\pi_n, \dots\}$. In (2), the connection between (2b) and (2a) is not explicitly marked, so the axioms of SDRT only yield a very underspecified contribution:

$$(13) \quad \text{sel}_{\mathbb{R}}(\pi_1, \text{sel}_{\Pi}\{\pi_2, \pi_3\}).$$

In words this means that some discourse relation selected from the set \mathbb{R} of discourse relations must hold between the constituent π_1 and one of the constituents introduced in (2b). Which discourse relation and which constituent are ultimately chosen to flesh out an underspecified formula depends on semantic and syntactic constraints and the surrounding discourse context. Combining the discourse units $\pi_1 - \pi_3$ from (2), their associated contents, (13), and $\text{Attribution}(\pi_2, \pi_3)$, we arrive at the logical form in (14). (14) is given in the SDRT language using a richly typed, dynamic, compositional framework (Asher and Pogodalla, 2011).

$$(14) \quad \exists \pi_1, \pi_2, \pi_3 (\pi_1 : \neg \text{come}(\text{john}, \text{my-party}) \wedge \pi_2 : \text{say}(\text{jill}, \pi_3) \wedge \pi_3 : \text{out-of-town}(\text{john}) \\ \wedge \lambda x \lambda y \text{sel}_{\mathbb{R}}(y, x)(\pi_1)(\text{sel}_{\Pi}\{\pi_2, \pi_3\}) \wedge \text{Attribution}(\pi_2, \pi_3))$$

⁴Were (2) embedded in a larger discourse, both π_2 and π_3 might be related to the preceding discourse. In fact, adding discourse connections can increase the unity and coherence of a discourse. This is captured in the principle Maximise Discourse Coherence from Asher and Lascarides (2003).

Let's now look at how we would resolve the underspecification in (14) in the context of a coercion story. The contents of π_1 and π_3 support $\text{Explanation}(\pi_1, \pi_3)$, and in fact, this is the most probable connection to π_1 in this case. Normally we would infer this relation to resolve the choice of rhetorical relation in (13), but because of the Attribution relation between π_2 and π_3 , we run into a conflict.

To make this conflict precise, we begin by assigning types to the variables π_1, π_2, \dots that represent discourse units in SDRT. There are two types for discourse representation variables, UP (U), for *uncommitted*, and DOWN (D), for *determinate*. The general idea is that the content associated with a variable of type D will be entailed by the discourse; for a variable of type U, neither its associated content nor the negation of this content will be entailed. Veridical relations, including Explanation, Elaboration, Narration, Contrast, and so on, are of type $D \rightarrow D \rightarrow D$. Non-veridical relations, including Conditional and Disjunction, have the type $U \rightarrow U \rightarrow D$. Attribution is a *right*-nonveridical relation, meaning that it is non-veridical in its right argument; that is, it has the type $D \rightarrow U \rightarrow D$.

The proposed typing for Attribution is designed to reflect the fact that the act of attributing content to an agent should not in itself require speaker commitment to the attributed content. As it stands, however, it does not capture the fact that the content inside of an Attribution's scope might itself contain relations that are veridical. For instance, consider:

(15) [Jill said that] $_{\pi_1}$ [John went running at 6] $_{\pi_2}$ [and he went to dinner with Julie at 7.] $_{\pi_3}$

where π_1, π_2 , and π_3 label the elementary discourse units.⁵ The SDRS K for (15) contains a complex discourse unit, π_0 , that itself contains a discourse relation. (16) gives its logical form.

(16) $\exists \pi_0, \pi_1, \pi_2, \pi_3 (\text{Attribution}(\pi_1, \pi_0) \wedge \pi_0 : \text{Narration}(\pi_2, \pi_3))$

The formula $\text{Attribution}(\pi_1, \pi_0)$ requires π_0 to be of type U, but this conflicts with the type D assigned to the veridical relation $\text{Narration}(\pi_2, \pi_3)$. A similar problem arises with conditionals:

(17) [If John doesn't have to work this evening] $_{\pi_1}$ [then he will go running] $_{\pi_2}$ [or he'll go play squash.] $_{\pi_3}$

We get the following logical form:

(18) $\exists \pi_0, \pi_1, \pi_2, \pi_3 (\text{Conditional}(\pi_1, \pi_0) \wedge \pi_0 : \text{Disjunction}(\pi_2, \pi_3))$

The relation Conditional requires π_0 to be of type U; Disjunction requires it to be of type D.

In the case of both Conditional and Attribution, we need to capture the fact that veridical relations in their scope are only veridical *relative to a context* introduced by the antecedent of the conditional or the attribution predicate, respectively. To do this, we relativize types to con-

⁵We use square brackets to mark the boundaries of discourse units.

texts.⁶ A context will be represented as a *segmented discourse representation structure* (SDRS) or formula under the scope of an operator; the outermost SDRS K_0 is a context, as is any SDRS K_i that is a constituent of K_0 and is within the scope of an operator. Note that a formula under the scope of a modal operator also serves as a typing context. If a discourse unit is typed D in a context K , which we will write as $\pi^K : D$, then its normal entailments are entailed in K ; more precisely, if $\pi^K : D$, the content associated with π must be satisfied at worlds or points of evaluation in which any relations in K involving π are satisfied. If a discourse unit π is such that $\pi^K : U$, then the content of π is not entailed in K even if it figures in a relation R that is entailed in K . By relativizing types to contexts, we capture the idea that while a speaker may not be committed to the content of a report, when she attributes content to another agent, she commits *to the other agent being committed* to the report's content. Note that in an SDRS, an instance of a discourse relation R occurs in one and only one context.

The relation between the entailments of a context K and the entailments of a context K' embedded in K obeys the following retyping rule: where K is a context, $\pi^K : u$, and where K' is a constituent of K under the scope of a modal operator or intensional relation, then $\pi^{K'} : d$ is an allowable typing. We also allow type raising from D to U when we shift from a context K' to a context K in which K' is embedded: if $\lambda \pi \psi$ requires a discourse unit of type U as argument in K , but is given an argument π_1 of type D from K' , then a general coercion permits us to rewrite the type of π_1 to U ; i.e. $\pi_1^K : U$. Finally, for any constituent π in K introduced by an indicative clause not in the scope of an intensional operator or intensional relation in K , $\pi^K : D$.

We say that an SDRS K has a *consistent type assignment* just in case: K respects the retyping rule; the discourse relations in K have arguments of the type required by the semantics of the relations; and for every context K' , and for every discourse unit π in K' , π has a unique type assignment relative to K' . The SDRS K for (16), for instance, has a consistent type assignment in which $\pi_1 : D$ and, after type raising, $\pi_0 : U$ in K . But π_0 is an SDRS in the scope of an intensional relation; it therefore constitutes a typing context of its own and specifies its own assignments to discourse units figuring in discourse relations in π_0 . In this context, π_0 is of type D —i.e., $\pi_0^{\pi_0} : D$; and we also have $\pi_1^{\pi_0} : D$ and $\pi_2^{\pi_0} : D$, which yields a consistent typing of K .

We note that examples like (19) do not require any kind of retyping or context shifting. While (19b) and the embedded clause of (19a) have the same propositional content, they contribute different discourse units and can therefore receive different type assignments.

- (19) a. Jill said John was out of town.
 b. And (indeed) he was out of town.

The types we have assigned to discourse units are distinct from the types assigned to contents, just as the contents associated with discourse units are distinct from the units themselves. The fact that (19b) serves as a correction or amplification of (19a) is because the same content is now linked to a label that commits the speaker to John's being out of town.

⁶For more on typing within contexts, see the notion of *coercive sub-typing*, e.g. Luo (1999).

The conflict that induces the coercion observed with the parenthetical report in (2b) is triggered by the fact that the Attribution relation and the Explanation relation in (2) share the same second argument, namely π_3 , but Explanation types this argument as D , while Attribution types it as U . More generally, discourse parenthetical reports will always trigger this kind of conflict; as explained in Hunter (2016), we do not find the embedded clauses of discourse parenthetical reports in non-veridical relations such as Disjunction or Conditional. The only available reading of (20), for example, is one in which *Linda said* takes scope over the entire conditional.

(20) If John finishes his housework, then Linda said he'll come to the party.

In this case, the report is not parenthetical.

To resolve the conflicting type requirements imposed by a formula of the form $\text{Attribution}(\pi, \pi')$ on the one hand and a veridical relation that takes π' as an argument on the other, something has to give: either the type assignment for a constituent, the type of the predicates over the constituents (i.e. the discourse relations), or the way these predicates combine with their arguments has to change. Our retyping rule will not solve the conflict: suppose we replace $\lambda x \lambda y \text{sel}_R(y, x)(\pi_1)(\text{sel}_{\Pi}\{\pi_2, \pi_3\})$ in (14) with $\text{Explanation}(\pi_1, \pi_3)$, which requires π_3 to have the type D . Raising π_3 from D to U will leave Explanation with an argument of type U , which conflicts with its veridical semantics. The retyping rule would only work if the entire Explanation were under the scope of the Attribution.

Retyping π_3 from U to D in $\text{Attribution}(\pi_2, \pi_3)$ will not work either. An important principle of the account of coercion in Asher (2011) is that lexical items or grammatically determined types do not change their type due to coercion, and Asher (2011) adduces considerable evidence for this principle. We believe that this principle is also plausible at the discourse level. Given this assumption and the fact that the Attribution relation is grammatically determined, we cannot change the type requirements of Attribution. It follows that π_3 will retain its typing of U in the context of the Attribution, and there is therefore no consistent type assignment for an SDRS that includes those two relations on π_3 .

Shifting the type of π_1 or π_2 would also be ineffective, because doing so would only introduce a conflict between these discourse referents and the requirements of Explanation and Attribution, respectively. Shifting the inferred link between π_2 and π_3 is not allowed because the link is grammatically determined. The only possibility, then, is to shift the inferred link between π_3 and π_1 . To spell this out formally, we need two things. We need to specify how the predication within a given context changes using a functor that takes a relation R and its two arguments α and β and returns a modified predication of R applied to α and β ; this is easy to do with our types and the availability of a modal operator. The other task is to specify what predicate licenses the introduction of the functor.

We begin with the second task. In standard coercion cases, the licensing predicate is typically R itself. For discourse coercions, however, it is the discourse environment that typically licenses the coercion. In particular, it is the presence of the Attribution that licenses the coercion of the environment involving the Explanation relation. Given the contents of π_1, π_2

and π_3 , the discourse model selects the most plausible specification for $\text{sel}_{\mathbb{R}}$ and sel_{Π} in (14): $\text{Explanation}(\pi_1, \pi_3)$. From the discourse model, we then rewrite the logical form in (14) as:

$$(21) \quad \exists \pi_1, \pi_2 (\pi : \neg \text{come}(\text{john}, \text{my-party}) \wedge \pi_2 : \exists \pi_3 \text{say}(\text{jill}, \pi_3) \wedge \pi_3 : \text{out-of-town}(\text{john}) \\ \wedge \lambda x \lambda y \text{Explanation}(y, x)(\pi_1)(\pi_3) \wedge \text{Attribution}(\pi_2, \pi_3))$$

But now we have a type conflict on π_3 , and λ reduction cannot proceed further: the type that Explanation requires of π_3 (D) conflicts with the type assigned to it by Attribution (U) in $\mathbf{K}_{(21)}$.

This type clash leads to the introduction of a coercion functor that can be specified as follows.

- Coercion to $\diamond R$: $f : \lambda \beta : U. \lambda \alpha : D. \lambda R : D \rightarrow D \rightarrow D. \diamond R(\alpha, \beta)$.

The presence of the Attribution entailing the typing $\pi_3 : U$ together with the simple assertion of (2a) entailing $\pi_1 : D$ and the discourse model's prediction of a veridical relation's holding of π_1 and π_3 license a rewriting of

$$(22) \quad \lambda x \lambda y \text{Explanation}(y, x)(\pi_1)(\pi_3)$$

as

$$(23) \quad f(\lambda x \lambda y \text{Explanation}(y, x))(\pi_1)(\pi_3)$$

λ -reduction can proceed now that the types all match, yielding:

$$(24) \quad \diamond \text{Explanation}(\pi_1, \pi_3)$$

So our final logical form for (2) looks like this:

$$(25) \quad \exists \pi_1, \pi_2 (\pi_1 : \neg \text{come}(\text{john}, \text{my-party}) \wedge \pi_2 : \exists \pi_3 \text{say}(\text{jill}, \pi_3) \wedge \pi_3 : \text{out-of-town}(\text{john}) \\ \wedge \diamond \text{Explanation}(\pi_1, \pi_3) \wedge \text{Attribution}(\pi_2, \pi_3))$$

Note that f is a conservative functor that maintains consistency with the original requirement of a veridical relation R such that $R \models \diamond R$. Note also, however, that \diamond itself specifies a typing context, so that within the context of \diamond it is consistent to have $\pi_1 : D$ and $\pi_3 : D$, as would be needed if a veridical relation had fallen in the scope of *Jill said* in (2b) (cf. (15)).

The functor f does not change the type of either α or β ; it rather shifts the relation inferred between them to a relation that is compatible with arguments of either type D or type U. Because the type of the arguments stays constant, the entailment in (11), repeated here as (26), does not reflect the actual entailments of an example like (2).

$$(26) \quad \diamond R(\alpha, \beta) \models \diamond(\alpha \wedge \beta)$$

In (2), the content of (2a) is asserted independently of the subsequent discourse, and is therefore

of type D; the functor simply renders the relation compatible with the typing of π_3 , which is U. We can therefore make the stronger claim for an example like (2):

$$(27) \quad \diamond R(\alpha, \beta) \models \alpha \wedge \diamond \beta$$

This conforms to the intuition that (2a) is entailed in (2).

This entailment requires some unpacking. The formula (25) entails, as described in (27), that the content of π_1 is satisfied at the actual world and the content of π_3 is satisfied at a world consistent with the speaker's commitments. Yet to the extent that a causal relation might hold between the eventualities e_α and e_β described by α and β , respectively, the Explanation relation as well as e_α and e_β must all hold in the same world. This is ensured by a fact about epistemic possibilities: given that π_1 is asserted, the speaker is committed to it, and so the content of π_1 holds at *all* worlds consistent with her commitments. What $\diamond \text{Explanation}(\pi_1, \pi_3)$ adds is the claim that in some of those commitment worlds, π_3 is also verified, and furthermore, in at least some of those worlds, π_3 describes an eventuality that caused that described by π_1 . In other words, under the scope of \diamond , π_1 and π_3 are of both type D, so the constraints on Explanation are satisfied.

Interestingly, there is an asymmetry in coercion facts. While (2) seems to have a discourse parenthetical reading of the sort we have derived in (25), an example that inverts the form of (2) does not.

- (28) a. Jill said John is sick.
 b. He ate a bad clam.

The claim in (28b) is naturally interpreted either as extending the report context introduced by (28a)—thus committing Jill to the content of (28b)—or as providing an explanation not only of why John is sick but also, in effect, of why Jill said what she did.

As defined, the functor f only works if the *second* argument of the coerced relation is U. In (28) it is the first argument of the implicit causal relation that is typed U, so our functor f will not apply to this case. Our account thus predicts that we can never have an instance of a veridical relation whose left argument is the embedded clause of a report and whose right argument is outside of the scope of the report (i.e. we can never coerce the argument of an Attribution to D).

This is surprising and may even seem suspicious. Yet some data, including (28), suggest that we should not countenance a functor g : $\lambda\beta: D. \lambda\alpha: U. \lambda R: D \rightarrow D \rightarrow D. \diamond R(\alpha, \beta)$. One reason for not stipulating the functor g is that alternative means of producing a modalized relation exist:

- (29) a. Jill said John is sick.
 b. Maybe it's because he ate a bad clam.

Simple assertions allow for modalisation in a way that reports do not. There is no way to

lexically induce the possible explanation reading for (2)—a speaker must either choose a report construction, which carries the extra information of a source, or a simple modal. Combining the two leads to non-parenthetical readings of the reports. So the presence of such alternatives for the inverted ordering might explain the asymmetry between our intuitions about (2) and (28).

On the other hand, there are examples, like (30), in which something like coercion of the first argument seems to work.

(30) [The school said] $_{\pi}$ [Isabel is sick.] $_{\pi_1}$ Maybe [she caught Rose's flu.] $_{\pi_2}$

This example implies $\diamond\text{Explanation}(\pi_1, \pi_2)$. However, note that (30) differs from (2) in that π_2 falls within the scope of a modal. We think that (30) manifests a kind of modal subordination involving the intensional context provided by $\text{Attribution}(\pi, \pi_1)$ and the modal operator contributed by *Maybe*. Modal subordination could allow for an attachment between the two discourse units of type U, yielding a formula of the form:

(31) $\exists \pi, \pi_1, \pi_2 (\text{Attribution}(\pi, \pi_1) \wedge \pi : \text{school said} \wedge \pi_1 : \text{sick(isabel)} \wedge \diamond(\text{Explanation}(\pi_1, \pi_2) \wedge \pi_2 : \text{caught(isabel, flu)}))$

Given that \diamond is a normal modality, (31) implies $\diamond\text{Explanation}(\pi_1, \pi_2)$.

(32) provides another case in which a $\diamond R$ is arguably inferred from a first argument of type U and a second argument of type D:

(32) [Where was John on Tuesday evening?] $_{\pi}$ [Jill said that] $_{\pi_1}$ [she and John went running at 6.] $_{\pi_2}$ [Then/After that, he and I had dinner at 7.] $_{\pi_3}$

(32) intuitively implies $\diamond\text{Narration}(\pi_2, \pi_3)$. In this case, however, it is the surrounding discourse structure that indirectly sets up the coercion on Narration. The answer to the question in π includes π_2 and π_3 . But let's concentrate first just on relating π , π_1 and π_2 . The connections between these units will parallel those between the units in (2), yielding $\diamond\text{QAP}(\pi, \pi_2)$, where *QAP* stands for *Question-Answer-Pair*. Now π_3 simply extends the answer that is being proposed by π_2 to a more complicated answer: perhaps the proposed sequence of events answers the question posed in π . Semantically, this means that in some world w from the set of worlds compatible with the speaker's commitments, π_2 and π_3 hold in w and the eventuality described by π_3 took place after the eventuality described by π_2 in w . In other words, there is an attachment of π_3 to π_2 forming a Narration under the scope of the modal introduced by the QAP, in a manner similar to the logical form for (16). This gives the following, abbreviated logical form:

(33) $\exists \pi, \pi_1, \pi_2, \pi_3, \pi_4. (\text{Attribution}(\pi_1, \pi_2) \wedge \diamond\text{QAP}(\pi, \pi_4) \wedge \pi : \text{Where was John?} \wedge \pi_1 : \text{Jill said} \wedge \pi_2 : \text{running} \wedge \pi_3 : \text{dinner} \wedge \pi_4 : \text{Narration}(\pi_2, \pi_3))$.

We note several consequences of (33). First, the Narration relation and its temporal consequences hold only between the events described in π_2 and π_3 , not the time of Jill's saying, as

intuitions dictate. Second, because of (27), $\diamond\text{QAP}(\pi, \pi_4)$ entails $\diamond\text{Narration}(\pi_2, \pi_3)$, as desired, even though this formula is not in the logical form. Third, this is the only possible SDRS given our constraints on consistent typing. We could not put $\text{Narration}(\pi_2, \pi_3)$ outside the scope of the $\diamond\text{QAP}$ relation in (33), because $\pi_2^{(33)} : U$ and this would conflict with typing demands of Narration as a veridical relation.

Example (32) does not provide a genuine motivation for adopting the functor g described above. $\diamond\text{Narration}$ is indeed entailed, but not because of a direct conflict between Narration and Attribution ; it's epistemic status follows from the fact that π_3 develops a context that is already modal. At the same time, given that the speaker asserts π_3 , its contents will hold at all worlds consistent with her commitments, including the actual world whether or not it is also a π_2 world. Thus we arrive at the following consistent top-level typing for (33): $\pi^{(33)} : D$, $\pi_1^{(33)} : D$, $\pi_2^{(33)} : U$, and $\pi_3^{(33)} : D$. Indeed, our typing rules mandate these. Our typing rules also permit $\pi_2^{\pi_4} : D$, $\pi_3^{\pi_4} : D$, and $\pi_4^{\pi_4} : D$ as part of a consistent typing.

Our account right now relies on the potentially controversial assumption that Attribution always types its second argument as U . We now turn to a discussion of this assumption in the context of factive embedding verbs.

4. Factives

On our proposal, the type assigned to a rhetorical relation R is entailed by R 's semantics; it is a property of the relation itself. As a result, the type is unaffected by semantic differences between connectives that license R . This means that Attribution , for example, will have the type $D \rightarrow U \rightarrow D$ even when triggered by a factive verb, as in (34).

- (34) a. John didn't come to my party.
 b. Jill found out that he's been working nights at the station.

This might at first seem counterintuitive. After all, a factive verb entails the truth of its complement, so shouldn't its complement have the type D ?

First of all, it is important to distinguish between the discourse contribution of a report and its interpretation. The discourse contribution determines what relation is added to the logical form for a discourse, but while these relations play a significant role in discourse interpretation, they do not tell the whole story. The content of discourse units certainly effects interpretation, and in the case of reports, world knowledge or opinions about the reliability of sources cited in reports (de Marneffe et al., 2012) can influence what information interpreters ultimately take away. There is therefore no problem with saying that discourse structure contributes certain entailments and that the content of discourse units, including the semantics of the various embedding verbs used in reports, can strengthen these entailments. A typing of U for a clause embedded under a factive verb is consistent with the demands of the verb, but the final discourse level entailments triggered by reports with factive verbs can ultimately be stronger. What we have modeled is the general function of discourse parenthetical reports, and we take as basic that the attribution predicate of such a report acts as a buffer between the embedded content and the

preceding discourse.

What's more, in many cases, the claim that discourse parenthetical reports involving factive embedding verbs trigger relations of the form $\diamond R$ is more than consistent with the demands of factive verbs—it is also borne out by the data. Even Attributions triggered by factive verbs allow a speaker to hedge her commitment to the discourse function of the embedded clause. If a speaker A asks, (a) “Why didn't John come to the party?”, a speaker B can reply, (b) “Jill found out that he has taken on a night job” and thereby hedge her commitment to the Question/Answer relation suggested by the contents of (a) and (b), even while committing to the truth of the embedded clause in (b) in virtue of the factivity of *find out*. Factivity licenses an inference *from* the discourse logical form but does not contribute type information to its *construction*, and thus doesn't affect the typing relevant to the interaction between attributions and the surrounding discourse context. This formalizes the observation in Simons (2007) that parentheticality and factivity are orthogonal notions.

Of course, intonation plays a role in the way that discourse parenthetical reports are ultimately interpreted. A more certain and declarative intonation for (b) above might encourage the interpretation of an Explanation rather than the weaker \diamond Explanation relation. Intonation certainly plays an important, if poorly understood, role in discourse interpretation. Nevertheless, even a declarative intonation does not *force* the stronger interpretation. Speaker B could continue (b) with (c): “That might be why he didn't come,” thereby making the \diamond Explanation explicit. Our point is that the use of a factive verb in a discourse parenthetical report does not in and of itself require a non-modal relation; the systematic discourse contribution of *all* reports is the weaker $\diamond R$. Strengthening of this relation at the level of discourse interpretation must be explained by independent factors.

5. Connectives

As we've seen, while the content of (2a) and the embedded content of (2b) together lead the discourse model to predict an Explanation relation, other information from (2) triggers a coercion. This coercion is blocked once we make the causal connection explicit as noted by Hunter and Danlos (2014). The report in (35b) can only receive a non-parenthetical reading.

- (35) a. John didn't come to my party
b. because Jill said he was out of town.

A consequence of our coercion story is that the typing demands of a connective with a grammatically determined scope must be satisfied. We predict this because the semantics of the connector *because* entails a non-modal commitment to Explanation over its complement clause, and our coercion account cannot change such grammatically given entailments.

The contrast between (36) and (37) also illustrates this point, though the data are more subtle.

- (36) a. Liz's album got a bad review.
b. John said she's really upset.

- (37) a. Liz's album got a bad review,
 b. so John said she's really upset.

In (36b), John's saying what he did might be completely independent of the review mentioned in (36a); (36b) arguably has a parenthetical reading and licenses \diamond Result. In (37b), to the extent that (37) is acceptable, we infer that John himself suspects that the review caused Liz's distress. The fact that John said what he did is thus a reaction to the review, so the matrix clause of (37b) describes a result of (37a). The matrix clause is of type D, so it satisfies the type demands of Result. At the same time, connecting (37b) to (37a) with Result is only really coherent if the embedded clause of (37b) is also related to (37a) with \diamond Result. But our coercion story still applies to any connection between the embedded clause of the report with (37a). So in fact we predict a graph with both Result and \diamond Result for (37).

Our prediction does not apply to all explicit connectives. Compare (38) with (35) or (39):

(38) [A man died] $_{\pi_1}$ [after police say/said] $_{\pi_2}$ [he was hit by a fire truck.] $_{\pi_3}$

(39) A man died because police said he was hit by a fire truck

While no inferred \diamond Explanation to the embedded discourse unit is possible in (35) or (39), an inference to \diamond Explanation is involved in the interpretation of (38) (Hunter and Danlos, 2014). This has to do with the semantics and discourse function of *after*, whose semantic scope is not determined by its syntactic position—in (38) the second argument of *after* is π_3 . While it marks a temporal relation or possible temporal relation in this case, a causal relation between π_1 and π_3 is also inferred, and there our coercion account applies.

Finally, we note a possible difference between syntactically embedded markers like *because*, connectives like *after* and markers that are marked with a comma, such as *after that* and *afterwards* (cf. Danlos, 2013). Consider the following data from (Hunter & Danlos 2014):

(40) John is very generous. For example, Jill said that he gave \$50 to a homeless man yesterday.

(41) John didn't come to the party. Instead, Jill said that he went to dinner with his brother.

Parenthetical readings for (40) and (41) are possible with explicit connectives, but once we embed the connective inside the report, we lose the parenthetical reading of (41):

(42) John didn't come to the party. Jill said instead that he went to dinner with his brother.

All of this suggests that while the typing demands imposed by discourse connectives must in general be respected, further work is needed to understand differences between the various types of discourse connectives.

6. Conclusion

Our main goal in this paper has been to develop a compositional account of the discourse function of parenthetical reports. Hunter (2016) argued that the role of discourse parenthetical reports is to contribute a modal discourse relation to a discourse structure. The derivation of such relations was, however, left unexplained and the \diamond relations that Hunter posited seemed more like primitive discourse relations. Such an approach does not take into account the fact that discourse parenthetical reports in fact arise from the interaction of an Attribution environment and the surrounding discourse context; the embedded content of a discourse parenthetical report and the discourse unit to which it attaches in the preceding discourse do not on their own support a relation of the form $\diamond R$. In the account developed here, $\diamond R$ is derived from R via coercion and a context sensitive notion of typing (Asher and Luo, 2012). The Attribution relation plays a key role in this story by introducing a typing on its second argument that is incompatible with the requirements of the relation R .

We have also extended our basic account to discourse parenthetical readings of factive verbs. In our account, all reports with non-parenthetical syntax are modelled with the discourse relation Attribution, and Attribution always has the type $D \rightarrow U \rightarrow D$, regardless of the embedding verb involved. We argued not only that this is consistent with the demands of factives and a reasonable understanding of how discourse structure contributes to discourse interpretation, but also that discourse parenthetical readings of even factive verbs can entail merely modal relations. As pointed out in Simons (2007), factivity and parentheticality are orthogonal notions.

Finally, we have shown how our account predicts that the attribution clause of a report in the scope of an explicit discourse connective must be relevant to that connective. Pure discourse parenthetical readings in which the attribution predicate is not related to the preceding discourse, or at least not related to the unit to which the embedded clause is related, are predicted to be blocked by explicit connectives because our coercion operation cannot shift grammatically determined types. We did note, however, that some connectives seem to allow parenthetical readings. A detailed case study of different connectives is needed to complete our account, but we think this account is promising enough to warrant such a study in the future.

References

- Asher, N. (1993). *Reference to Abstract Objects in Discourse*. Number 50 in Studies in Linguistics and Philosophy. Dordrecht: Kluwer.
- Asher, N. (2011). *Lexical Meaning in Context: Web of Words*. Cambridge University Press.
- Asher, N. and A. Lascarides (2003). *The Logics of Conversation*. Cambridge University Press.
- Asher, N. and Z. Luo (2012). Formalisation of coercions in lexical semantics. In *Sinn und Bedeutung*, Volume 17, pp. 63–80.
- Asher, N. and S. Pogodalla (2011). Sdr and continuation semantics. In T. Onada, D. Bekki, and E. McCready (Eds.), *New Frontiers in Artificial Intelligence: JSAI-isAI 2010 Workshops, LENLS, JURISIN, AMBN, ISS, Tokyo, Japan, November 18-19, 2010, Revised Selected Papers*, pp. 3–15. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Danlos, L. (2013). Connecteurs de discours adverbiaux: Problèmes à l’interface syntaxe-

- sémantique. *Linguisticae Investigationes* 36(2), 261–275.
- Danlos, L. and O. Rambow (2011). Discourse relations and propositional attitudes. In *Proceedings of the Constraints in Discourse Workshop (CID 2011)*, Agay, France.
- de Marneffe, M.-C., C. D. Manning, and C. Potts (2012). Did it happen? the pragmatic complexity of veridicality assessment. *Computational Linguistics* 38(2), 301–333.
- Dinesh, N., A. Lee, E. Miltsakaki, R. Prasad, and A. Joshi (2005). Attribution and the (non-)alignment of syntactic and discourse arguments of connectives. In *Proceedings of ACL Workshop on Frontiers in Corpus Annotation*, Ann Arbor, MI, USA.
- Hardt, D. (2013). A uniform syntax and discourse structure: the copenhagen dependency treebanks. *Discourse and Dialogue* 4(2).
- Hobbs, J. (1985). On the coherence and structure of discourse. Technical report, Center for the Study of Language and Information, Stanford University. Report CSLI-85-37.
- Hooper, J. B. (1975). On assertive predicates. In J. Kimball (Ed.), *Syntax and Semantics*, Volume 4, pp. 91–124. NY: Academic Press.
- Hunter, J. (2016). Reports in discourse. *Dialogue & Discourse* 7(4), 1–35.
- Hunter, J. and M. Abrusán (2016). Rhetorical relations and QUDs. In *New Frontiers in Artificial Intelligence: JSAI-isAI Workshops LENLS, JURISIN, KCSD, LLLL Revised Selected Papers*. Springer. forthcoming.
- Hunter, J., N. Asher, B. Reese, and P. Denis (2006). Evidentiality and intensionality: Two uses of reportative constructions in discourse. In *Proceedings of the Constraints in Discourse Workshop (CID 2006)*, Maynoth, Ireland.
- Hunter, J. and L. Danlos (2014, April). Because we say so. In *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL)*, Gothenburg, Sweden, pp. 1–9. Association for Computational Linguistics.
- Luo, Z. (1999). Coercive subtyping. *Journal of Logic and Computation* 9(1), 105–130.
- Mann, W. and S. Thompson (1988). Rhetorical structure theory: Towards a functional theory of text organization. *Text* 8, 243–281.
- Ross, J. (1973). Slifting. In M. Gross, M. Halle, and M. Schutzenberger (Eds.), *The formal analysis of natural languages*, The Hague, pp. 133–169. Mouton.
- Simons, M. (2007). Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117(6), 1034–1056.
- Urmson, J. O. (1952). Parenthetical verbs. *Mind* 61 (244), 480–496.