

Reports in Discourse

Julie Hunter

juliehunter@gmail.com

*IRIT, Université Paul Sabatier
Toulouse, France &
GLiF, Universitat Pompeu Fabra
Barcelona, Spain*

Editor: Jonathan Ginzburg

Submitted 08/2014; Accepted 06/2016; Published online 06/2016

Abstract

Attitude or speech reports in English with a non-parenthetical syntax sometimes give rise to interpretations in which the embedded clause, e.g., *John was out of town* in the report *Jill said that John was out of town*, seems to convey the main point of the utterance while the attribution predicate, e.g., *Jill said that*, merely plays an evidential or source-providing role (Urmson, 1952). Simons (2007) posits that parenthetical readings arise from the interaction between the report and the preceding discourse context, rather than from the syntax or semantics of the reports involved. To my knowledge, however, no account of these discourse interactions has been developed in formal semantics. Research on parenthetical reports within frameworks of rhetorical structure has yielded hypotheses about the discourse interactions of parenthetical reports, but these hypotheses are not semantically sound. The goal of this paper is to unify and extend work in semantics and discourse structure to develop a formal, discourse-based account of parenthetical reports that does not suffer the pitfalls faced by current proposals in rhetorical frameworks.¹

Keywords: Speech reports, Parenthetical reports, Evidential reports, Discourse structure, Discourse Connectives

1. Introduction

In the absence of further context, a speaker² who utters (1) would naturally be understood as using (1b) to propose an explanation for why John did not come to her party.

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

The content of *Jill said* in (1b) does not participate, at least not directly, in this explanation—while it is possible to imagine a scenario in which John didn't come to the party because *Jill said* he was out of town, this reading of (1b) would require a more elaborate discourse than (1). It is rather the

1. I would like to thank Márta Abrusán, Pascal Amsili, Nicholas Asher, Laurence Danlos, Eric Kow, Mandy Simons, participants of the *CREST International Workshop on Formal and Computational Semantics* at Kyoto, and participants at the Nijmegen workshop *Backgrounded Reports*, including organizers Corien Bary and Emar Maier, for helpful discussions on the issues discussed in this paper. I also thank my three anonymous *Dialogue & Discourse* reviewers for their thorough and constructive comments. This work has been supported by the French agency *Agence Nationale de la Recherche* (ANR-12-CORD-0004) and the *European Research Council* (grant 269427).

2. Throughout the paper, I will use *speaker* to refer to the agent of either a spoken or written utterance/sentence token.

embedded clause, *he was out of town* that is understood as a (possible) explanation of John's absence. Informally, we can say that the embedded clause seems to contribute the main point of (1b) while the attribution predicate, *Jill said*, merely serves to provide a source or evidence for the embedded content. Following Hooper (1975), Simons (2007) and Urmson (1952), I will call the use of *say* in (1b) a *parenthetical* use and I will call a report in which the embedding verb is used parenthetically a *parenthetical report*.

Now compare (1) with (2), which involves a standard, non-parenthetical use of the same report.

- (2)
- a. John is mad at Jill.
 - b. Jill said he was out of town,
 - c. so I didn't invite him to my party.

The report in (2b) is identical to that in (1b), yet the report does not make the same contribution to (2) as it does to (1). In (2b), the attribution predicate is integral to the explanation of (2a): the speaker is not suggesting that John is mad at Jill because he was out of town, but that he is mad at Jill *because she said something* (which led to his being excluded from the speaker's party). Thus, while the embedded clause alone seems to make the main discourse contribution of (1b), the attribution predicate of (2b), or perhaps the report as a whole, crucially participates in the report's main contribution to (2).

The discursive difference between (1b) and (2b) brings with it a difference in semantic entailments. To the extent that a speaker of (1) is committed to the possibility that John didn't come to the party because he was out of town, she must also be committed to the possibility that John was out of town (at the relevant time). That is, she must be committed to at least the possibility that the embedded content of (1b) is true. There is no such requirement of commitment to the embedded content of (2b): (2b) can be used to explain (2a) even in a context in which it is common knowledge that John was *not* out of town.

If we assume that parenthetical reports like (1b) have the same syntactic structure as their non-parenthetical counterparts—which I, following Simons (2007), will—then we must conclude that the discursive and semantic differences between (1b) and (2b) do not arise from the syntactic and semantic features of the report alone, but rather from the interaction between these features and other discourse moves. Accordingly, I will introduce the more specific term *discourse parenthetical reports* to talk about parenthetical reports akin to (1b). This will help to distinguish these reports from those whose parenthetical readings are marked syntactically (see §5 for further discussion of syntactic parentheticals).³

The goal of this paper is to develop a formal, discourse-based model of discourse parenthetical reports. While parenthetical reports have been discussed at length in Hooper (1975), Rooryck (2001), Simons (2007), and Urmson (1952) and it has been noted that parenthetical readings arise from the discourse function of parenthetical reports, there is as of yet no formal model of their discourse function or of how this function affects the semantic entailments of reports in different discourse contexts. Work in formal semantics has tended to bring observations about parenthetical readings back to bear on outstanding problems in formal semantics. Urmson, for example, was concerned with the implications of parenthetical readings for the semantics of attitude reports. Simons uses parenthetical reports to argue that the presuppositional behavior of factive verbs is not determined by their lexical semantics: given the right discourse context, even the content in the scope of a factive verb can carry the main point of the report—i.e., the report can have a parenthetical reading—and so the

3. I would like to thank an anonymous reviewer for suggesting the term *discourse parenthetical*.

embedded content can fail to be presupposed. This paper takes for granted the syntax and semantics of reports, and in particular, the assumption that reports with the surface structure of (1b)/(2b) have the same semantics and syntax regardless of whether they are interpreted parenthetically or not. The aim is rather to formalize the notion of discourse function at work in parenthetical readings and to explain how it accounts for the different entailments that arise from parenthetical and non-parenthetical readings.

Simons provides the most developed discussion of the discourse function of parenthetical reports that I am aware of in the formal semantics and pragmatics literature, but as her focus is on the behavior of different embedding verbs, she limits her discussion to how reports behave in question/answer sequences:

- (3) A: Why didn't John come to my party?
 B: Jill said (thinks, suspects, imagines, supposes, heard...) that he's out of town.

In these sequences, Simons proposes that the discourse function of the embedded clause of a discourse parenthetical report, e.g. (3B), is to answer the question posed by the preceding move, e.g. (3A). She says: "whatever proposition communicated by the response constitutes an answer (complete or partial) to the question is the main point of the response" (p. 1036). As Simons explicitly acknowledges (p. 1035), however, this criterion provides only a limited picture of the discourse function of parenthetical reports.

Theories of rhetorical structure such as *Rhetorical Structure Theory* (RST; Mann and Thompson (1988)) and *Segmented Discourse Representation Theory* (SDRT; Asher and Lascarides (2003)) have more developed notions of discourse structure and discourse function. The discourse function of an utterance u is given by the discourse relation that connects the content c of u to the surrounding discourse: if c is related to some other content c' via an Explanation relation, then c 's discourse function is to explain the eventuality described by c' ; if c is related to c' via a Narration relation, then its discourse function is to push the narrative forward by describing the next (discourse relevant) eventuality that occurs after that described by c' , and so on. The structure of the discourse is then determined by the collection of relation instances in the discourse.

This notion of discourse function has been applied to discourse parenthetical reports in multiple efforts to annotate newspaper texts for rhetorical structure (Dinesh et al., 2005; Hardt, 2013; Hunter et al., 2006). The parenthetical/non-parenthetical distinction generally becomes relevant when annotators, who are in many cases untrained in linguistics, decide that only the embedded content of a report is relevant to the discourse relation that connects the report to the preceding discourse. (4), an attested example with the form of (1), provides an illustration.

- (4) *London serves increasingly as a conduit for program trading of U.S. stocks. Market professionals said **London has several attractions**. First, the trading is done over the counter ... Second, it can be used to unwind positions before U.S. trading begins...* (PDTB, File 0097).

(4) is originally from *The Wall Street Journal*, but figures in the corpus for the Penn Discourse Tree Bank (PDTB; Prasad et al. (2007b)). In this example, PDTB annotators chose the content of the boldface text as the proposed explanation for the claim expressed by the sentence in italics; the attribution predicate of the report, *Market professionals said*, was excluded from the explanans. In other words, annotators took the report in (4) to be discourse parenthetical.

While discourse parenthetical reports might not be discussed as such in these annotation efforts, the annotation schemes that they have developed in response to discourse parenthetical readings in effect yield the following definition: a report r is discourse parenthetical just in case it is the embedded clause of r , not the attribution predicate, that enters into a discourse relation with some element from the discourse preceding r .⁴ Simons' proposed constraint for question/answer pairs can be seen as a particular case of this more general notion: a report r is discourse parenthetical just in case the embedded clause of r alone provides the second argument for an instance of the relation Question/Answer Pair, where q provides the first argument, for some question q in the incoming discourse.⁵ Of course, Simons' constraint could be generalized in different ways. An alternative approach might involve defining discourse function and structure within a Question Under Discussion account (Ginzburg, 2012; Roberts, 2012; Simons et al., 2010), though I will not develop such an approach here.

The model that I develop builds on the more general notion of discourse function offered by rhetorical theories, but rectifies certain semantic problems inherent in the annotation-driven solutions that have been offered. In particular, as I explain in §3, extant proposals do not account for the fact, well-known from semantic discussions of parenthetical reports, that a speaker who utters a discourse parenthetical report need not in general be fully committed to the embedded content of that report. Nor do they account for the fact that a speaker must nevertheless be at least tentatively committed to the embedded content. Finally, these accounts prevent the attribution predicate and the embedded clause of a report from being simultaneously relevant to the discourse outside of the report, at least when they play distinct rhetorical functions. These accounts in effect require annotators to treat either one or the other clause as the "main point" of the report, but that requirement is too strong given the data on reports in discourse. In a nutshell, my proposal is that a discourse parenthetical report is one in which the embedded clause is rhetorically connected to the discourse preceding the report, as in extant accounts, but unlike extant accounts, the embedded clause is discursively and semantically subordinate to the attribution predicate. As a result of these rhetorical connections, discourse parenthetical reports introduce modal discourse relations between the embedded clause and the preceding discourse, and the attribution predicate can be rhetorically related to the preceding discourse, although it must be related by a different discourse relation.

It should be noted that the account that I offer is designed to model the contribution of discourse parenthetical reports to discourse structure, but is not intended as a complete semantic account of parenthetical reports. It is meant to complement, rather than replace, studies on other factors that affect the interpretation of speech or attitude reports such as the lexical semantics of embedding verbs, e.g. Simons (2007), or the influence of certain kinds of world knowledge on our judgements about the reliability of a discourse parenthetical report, e.g. de Marneffe et al. (2012). It should also be a useful supplement to annotation-based work on rhetorical structure. Determining the impact of speech and attitude reports on discourse structure and interpretation is of utmost importance. For tasks such as automated discourse parsing, text summarization, and the automatic recognition of textual entailment, for example, we need to be able to draw reliable inferences from a discourse as

4. Constraints on how a bit of content can attach to the incoming discourse context are given by independent principles that vary from theory to theory. All that is important here is that the embedded clause is being attached to a representation of the discourse prior to the report, rather than to its own attribution predicate.

5. Different frameworks might have different names for this relation; I've opted here for the terminology of SDRT, but that is not important. However, given the semantics of Question/Answer Pair (QAP) in SDRT (Asher and Lascarides, 2003), it is important to note that this criterion will be equivalent to Simons' because for a discourse unit to serve as the second argument to an instance of QAP, it must provide a complete or partial answer to the first argument.

a whole. An important piece in this puzzle is figuring out how speakers make use of other peoples' speech acts and attitudes to perform their own speech acts and convey their own attitudes.

I begin in §2 by providing more detail on rhetorical theories and reviewing three proposals for the annotation of discourse parenthetical reports within different rhetorical frameworks. While these accounts differ in various ways, they all share the core idea that discourse parenthetical reports should be modelled by attaching the embedded clause to the preceding discourse. In §3, I show that these accounts as they stand come into conflict with semantic facts about discourse parenthetical reports and with independent principles of rhetorical structure. In §4 I show how the notion of discourse function adopted in rhetorical theories can be developed into a more more general, consistent account of discourse parenthetical reports. §5 takes a look at how syntactic parentheticals fit into the picture developed in §4. §6 concludes the discussion.

2. Rhetorical theories

Extant rhetorical theories, including RST, SDRT, D-LTAG (Webber, 2004), and other annotation methods, including those for the PDTB and the Copenhagen Discourse Tree Bank (Buch-Kromann and Korzen, 2010), differ in critical ways, and in the more theoretical discussion below, we will need to focus on a single theory. Nevertheless, the import of the current study should be of interest for research on rhetorical structure more generally. While the three accounts that I describe below are worked out using very different annotation methodologies, the heart of the proposals is the same, and as such, parts of the discussion to follow will be relevant to all of them. Where we will need a particular theory is in working out the details of my concerns and developing a positive proposal.

Building a representation of the rhetorical structure of a given discourse requires performing three tasks: segmenting the discourse into minimal units, attaching each discourse unit to some other unit in the structure, and labelling each discourse attachment with a rhetorical relation. When it comes to the annotation of discourse parenthetical reports, the approaches outlined in the PDTB (Dinesh et al., 2005), the Copenhagen Dependency Treebank (CDT) (Buch-Kromann et al., 2011), and SDRT (Hunter et al., 2006; Reese et al., 2007) all follow a similar recipe for accomplishing these tasks. First, a report r that is intuitively discourse parenthetical is segmented such that the embedded clause contributes its own discourse unit u . Second, u is attached to another unit u' in the discourse structure s built from the set of utterances prior to that of r . Third, the relation between u' and u is labelled with the discourse relation that intuitively would have related u' and u had u not been the argument of an attribution predicate. Finally, if desired, the discourse function of the attribution predicate can be modelled by attaching it to u with a special discourse relation that indicates its evidential (or emotive, etc.) discourse function.

The PDTB, for example, would indicate the discourse contribution of (1) as follows:

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

Following the conventions of the PDTB manual (Prasad et al., 2007b), content that contributes to the first argument of a discourse connective is placed in italics, while content that contributes to the second argument is placed in boldface. The connective, when explicit, is underlined; implicit connectives are marked as in (1b). We can see from the representation of (1) (which echoes the actual PDTB annotation of the found example (4)) that the embedded clause forms its own segment and it alone serves as the argument that ties the report to the discourse context prior to (1b) (Prasad et al.,

2007a).⁶ This segmentation choice contrasts with the segmentation for non-parenthetical reports, in which the entire report, attribution predicate plus embedded clause, is treated as a single discourse segment. The PDTB does not take a stand on the discourse contribution of the attribution predicate in discourse parenthetical reports; information about the sources and spans of attributions is stored alongside the annotation of a discourse in the PDTB but there is no connective posited to capture the discourse function of this information.

In the CDT, the annotation of syntactic structure and discourse structure is done using a single dependency graph so that there is a high level of uniformity between discourse and syntactic structure (Buch-Kromann and Korzen, 2010). One of the rare violations of this uniformity is allowed for discourse parenthetical reports, in which only the embedded clause is treated as a discourse argument. Instances of relations involving discourse parenthetical reports (and other syntax/discourse mismatches) are marked with a ‘*’: an asterisk to the left of a discourse connective signals that the discourse/syntax mismatch can be found in the first argument of the relation; an asterisk to the right signals that the mismatch lies in the second argument (Hardt, 2013). (1), for example, would be annotated along the following lines: Explanation*(1a, 1b). Like the PDTB, the CDT does not take a stand on the discourse contribution of the attribution predicate in a discourse parenthetical report.

A third proposal was offered by Hunter et al. (2006) for SDRT as part of the project DiSCoR⁷. Hunter et al., in contrast to the PDTB group, always treat the attribution predicate and embedded clause of a report as separate segments, regardless of whether the report has a discourse parenthetical reading or a non-parenthetical one. The different readings of reports are distinguished by the fact that the embedded clause of a discourse parenthetical report is the clause that is attached to the incoming discourse structure, as for the PDTB and CDT. In addition, discourse parenthetical and non-parenthetical reports are distinguished by different discourse relations that connect the attribution predicate of the report to the embedded clause: Attribution is used for non-parenthetical reports and Source, for parenthetical reports.⁸ In Attribution, the embedded clause is discourse subordinate (see Asher and Lascarides (2003)) to the attribution predicate, echoing the syntactic structure of the report. Source, however, reverses the arguments of Attribution so that the attribution predicate is discourse subordinate to the syntactically embedded clause. Furthermore, Source and Attribution differ in their entailments: when two arguments are related by Source, the text entails the content of both arguments, whereas when two contents are related by Attribution, the text only entails that the agent of the attribution stand in the relation indicated by the embedding verb to the content of

-
6. The PDTB group argue that the relation of attribution is one that holds between an individual and an abstract object. As discourse relations are taken to relate abstract objects only, Attribution is not treated as a discourse relation. So the difference between discourse parenthetical reports and non-parenthetical reports is that in the latter, the whole report is treated as a single unit that can figure in an argument for a discourse relation whereas in the former, only the embedded content contributes to an argument. In their words: “a discourse relation may hold either between the attributions (and the agents of attributions) themselves or only between the abstract object arguments of the attribution...” (p. 20)
7. DiSCoR, *Discourse Structure and Co-reference Resolution*, was an NSF project designed to study the relation between co-reference resolution and discourse structure. Most of the texts annotated for discourse structure were texts from the Message Understanding Conference (MUC) 6, that were already annotated for co-reference.
8. The relation Source was originally introduced by Hunter et al. (2006) under the name *Evidence*. The name was changed so that *Evidence* could be used for a different evidential relation. I have chosen to use the name *Source* here because it is consistent with later SDRT annotations, e.g. Reese et al. (2007), that incorporated Hunter et al.’s approach.

the embedded clause.⁹ An example of Source is provided in (1H), following the segmentation of (1) below:

- (1) a. [John didn't come to my party.]_α
 b. [Jill said]_β [he was out of town.]_γ
- (1H) Explanation(α, γ), Source(γ, β)

Two further accounts of speech reports that deserve mention here, though I will not discuss them in the rest of the paper, are those of Carlson and Marcu (2001) and Redeker and Egg (2006), both developed in *Rhetorical Structure Theory* (Mann and Thompson, 1988). Carlson and Marcu suggest that all reports be annotated with a relation that they call *Attribution*, which is structurally similar to Hunter et al.'s Source in that the attribution predicate in a report is treated as the satellite and the reported content serves as the nucleus. Redeker and Egg (2006) criticizes Carlson and Marcu's approach and offers an account that reverses the arguments of Attribution so that the reported content is the satellite and the attribution predicate, the nucleus. The reason why I will not pursue either of these accounts in this paper is that neither makes a distinction between discourse parenthetical reports and non-parenthetical reports.¹⁰ The attribution predicate is simply treated as the satellite of Attribution in Carlson and Marcu (2001) and as the nucleus in Redeker and Egg (2006). What we're interested in for the purposes of this paper is accounts that advocate a solution specifically for discourse parenthetical readings of reports.

The treatment of discourse parenthetical reports in the PDTB, the CDT, and Hunter et al. (2006) all have in common the idea that discourse parenthetical reports are best modelled by attaching the embedded clause directly to the incoming discourse and that this attachment pattern distinguishes them from non-parenthetical reports, in which it is the attribution predicate that is attached to the incoming discourse. Accordingly, I classify these accounts as *attachment* solutions. However, the claim that discourse parenthetical reports show different attachment patterns in discourse does not entail that discourse parenthetical reports and non-parenthetical reports should be distinguished syntactically, and I will take it for granted that the difference is not a syntactic one—except, of course, when the parenthetical verb appears in a syntactic parenthetical as in *Mary will be late, John said* (see §5 for a discussion of syntactic parentheticals). I will not defend this position here, because my focus is on the discourse contribution of parenthetical reports, regardless of their syntactic structure.¹¹ The point is that claims about the discourse structure of a chunk of discourse do not automatically entail claims about the syntactic structure of the constituents in that chunk. In fact, of the frameworks introduced here, the one that assumes the closest tie between syntactic structure and discourse structure is that for the CDT, and even the CDT treats discourse parenthetical reports as involving a syntax/discourse *mismatch*. In other words, they assume that the discourse contribution of a parenthetical report does not mirror its syntactic form.

9. For Attribution(α, β) or Source(β, α), the content of α will entail: *x said (thought,...) p* for some agent *x*, and β will specify the content of *p* in the sense that *p* will denote a subset of the worlds in the proposition denoted by K_β , i.e. the DRS that results from the processing of β .

10. However, see Matthiessen and Thompson (1988) for an interesting discussion about the fact that whether a given clause is ultimately considered to be a main clause or a subordinate clause depends on features of the discourse in which that clause is used.

11. See Simons (2007) for arguments that reports in which an embedding verb appears in the main clause, as in *John said (that) Mary will be late*, have the same syntactic structure regardless of whether the embedding verb is used parenthetically or not.

3. Conflicting criteria

In the ensuing discussion, I will adopt SDRT as my theoretical framework. To fully model the rhetorical contribution of reports, we need to take a stand on the discourse function of both the embedded clause and the attribution predicate, and we need a theory that offers independent principles to guide our theoretical choices on this matter. SDRT satisfies this criterion, but not all rhetorical frameworks do. The PDTB annotation method, for instance, is designed to be theory-neutral, and so cannot provide a theoretical framework by design. Moreover, part of its theory-neutral approach is to annotate isolated pairs of discourse arguments and the connectives that relate them; there is no goal to describe the discourse contribution of every discourse unit. As a result, the PDTB can remain agnostic about the role of the attribution predicate in discourse parenthetical reports.

SDRT is also largely motivated by semantic and pragmatic concerns and is the only rhetorical theory to provide a semantics for its discourse relations and to deliver fully interpretable logical forms for discourse. Many of the issues that we confront with discourse parenthetical reports are semantic and pragmatic, having to do with tracking the entailments of reports in a discourse or exploring aspects of their behavior that cannot be traced back to their syntactic structure. Thus a theory like SDRT is preferable to the framework of the CDT, which is wedded to a strong correspondence between syntactic structure and discourse structure, and even to RST, which does not emphasize the semantic interpretation of its discourse structures.

I will also limit the range of report verbs in my study. There are many factors that influence the interpretation of reports aside from rhetorical structure: lexical semantics, world knowledge, perhaps focus, and so on. To study the interaction of rhetorical structure and reports, we need to minimize the influence of these other factors as much as possible. Most of my discussion will therefore be centered around embedding verbs like *say* and other speech report verbs that are non-factive and so do not indicate a particular level of author commitment to the content in their syntactic scope. Such speech report verbs are common in the corpora that I am pulling from¹² and easily give rise to both discourse parenthetical and non-parenthetical readings. Also common are third person reports, and my discussion will focus largely on these as well, as it is with third person reports that questions about author commitment and entailments of reported content really become tricky.

With these caveats in place, I turn now to an argument that the attachment-based solution outlined in section 2 leads to conflicts with semantic facts about discourse parenthetical reports and with independent principles of rhetorical theories.

3.1 Hedged commitments

A speaker will often use a discourse parenthetical report to weaken or “hedge” her commitment to the embedded content of the report (Simons, 2007). In (1), for example, if the speaker were sure that John had been out of town, then the simplest solution would be to say so directly; the fact that she does not suggests that she is not entirely certain that he was out of town. Of course, we can imagine scenarios in which the use of a discourse parenthetical report is compatible with full speaker commitment to the embedded content, but what’s important is that a discourse parenthetical report does not itself require such commitment.

This fact about discourse parenthetical reports comes into conflict with the constraint of *veridicality* imposed by many discourse relations. A relation *R* is *veridical* just in case the truth of an

12. The majority of the verbs in the corpus annotated for DisCoR were evidential verbs, e.g. *say*, as opposed to, for example, verbs indicating emotions, e.g. *regret*.

instance of R entails the truth of that instance's arguments. Boolean conjunction, for example, is veridical: a formula of the form $p \wedge q$ can only be true in a model M if both p and q are true in M . Conditional relations, by contrast, are not veridical: a formula of the form $p \rightarrow q$ can be true even if both of its arguments are false. SDRT incorporates these basic logical connections into the semantics of its discourse relations—an arguably reasonable move for any semantic theory of rhetorical structure. The relation Continuation, for instance, conjoins two discourse units, and so each instance of Continuation will entail both of its arguments. Contrast also has conjunction as the foundation of its semantics, as does Narration; these relations are therefore veridical as well. Explanation is yet another veridical relation: a discourse unit u cannot truly explain another discourse unit u' unless u' and u both describe eventualities that held in the world of evaluation. To capture this dependence, a discourse formula of the form $Explanation(u', u)$ can only be true in a model M in SDRT if u and u' are true in M .

When a speaker performs a speech act that presents a discourse unit u as standing in a relation R to another unit u' , the content of this speech act, $R(u', u)$, is added to the logical form for the discourse (for any R, u, u'). If R is veridical, then the speaker takes on a commitment to the content of both u and u' , i.e. the content of her discourse entails both u and u' .¹³ Note that this does not ensure that either u or u' will be *true* in the relevant model—speakers are not infallible—it only ensures the speaker's *commitment* to their truth. Here is where the conflict between veridicality and *hedged commitments* arises: attachment solutions attempt to model the discourse function of discourse parenthetical reports by attaching the embedded content of a report to the incoming discourse with the same relation that would have been used had the embedded content not been embedded. Where the relation is veridical, this entails author commitment to the embedded content—commitment that the speaker may not be ready to take on. What's more, the relation at issue always *will be* veridical. If a discourse parenthetical report seems to provide an argument to a non-veridical relation, such as Alternation or Conditional, the result is that the attribution predicate will be understood as scoping over the discourse relation, as in (5):

(5) If [John finishes his housework] $_{\alpha}$, then [Linda said] $_{\beta}$ [he'll come to the party.] $_{\gamma}$

If Linda has reasons for thinking that John will come to the party that have nothing to do with him finishing his housework, then the choice of antecedent in (5) is unmotivated. The natural interpretation is therefore one in which Linda is committed to the conditional as a whole.

Hunter et al. (2006) avoid the conflict between veridicality and *hedged commitments* by requiring that Source be veridical. This means that Source can only be used in cases in which annotators (interpreters) judge that the speaker is committed to the embedded content. Recall Hunter et al.'s analysis of (1) repeated here:

- (1) a. [John didn't come to my party.] $_{\alpha}$
 b. [Jill said] $_{\beta}$ [he was out of town.] $_{\gamma}$

(1H) Explanation(α, γ), Source(γ, β)

In fact, while I provided this example to illustrate the structural features of *Source*, this annotation would have only been allowed by Hunter et al. if the context provided reason to accept Jill's report as totally reliable.

13. This picture of commitment follows original SDRT, but is complicated by issues such as embedded commitments and disagreements about commitments, topics which are handled in more recent work on SDRT. See Venant et al. (2014) and Venant and Asher (2016).

In the corpora used by the PDTB, the CDT, and DiSCoR, which consist mainly of newspaper articles, treating Source as a veridical relation often yields the right results. This is because in many of the articles, the content is so uncontentious and the sources so reliable that there is little reason to question the embedded content of the reports or the author's commitment to it. (6b) is discussed by Dinesh et al. (2005) as an example of Contrast with a discourse parenthetical report (their example (12)):

- (6)
- a. At the same time, the New Brunswick, N.J., company said negotiations about pricing and volumes of product had collapsed between it and its exclusive distributor in the U.S., National Medical Care Inc. ... Yesterday, the [Delmed] spokeswoman said sales of Delmed products through the exclusive arrangement with National Medical accounted for 87% of Delmed's 1988 sales of \$21.1 million.
 - b. *The current distribution arrangement ends in March 1990, although Delmed said **it will continue to provide some supplies of the peritoneal dialysis products to National Medical**, the spokeswoman said.*

(6a) provides two of the three sentences that precede (6b) in the PDTB (file 0970). (6b) is presented as being factual: Delmed and the spokeswoman should be reliable sources given their relation to Delmed, there is nothing contentious about the content that they have reported (the content that Dinesh et al. mark as contributing to Argument 2 of *although*), and the author gives no signal in the rest of the article that s/he is not fully committed to this content. In examples like (6b), the take-home message seems unchanged if we remove information about the writer's sources. (7) illustrates the same point, but with an Elaboration or Instance relation. While this example is not from the PDTB, I will use their conventions for marking the arguments for ease of exposition.

- (7)
- a. Another firm, uReveal, thinks that it's cracked the code. Charles "Bucky" Clarkson, uReveal's chairman and CEO, said that software such as his makes it easier to to parse all those government reports and organize the data so that analysts can get more out of it, and more quickly. He also claims that the software is so simple to use that (gasp!) even liberal arts majors can use it.
Jokes about "soft majors" aside, the idea of data analysis tools easy for anyone to use is compelling because it frees up data scientists to do more specialized work.
 - b. *It can also bring in specialist knowledge from people who aren't data scientists.* Clarkson said, for example, that **deploying these kinds of data analysis tools in hospitals have allowed doctors to spot trends that they would have otherwise missed**, by analyzing their observation notes in conjunction with other electronic medical records.

One doesn't have to look far, however, even in the realm of newspaper articles, to see that discourse parenthetical reports are used widely in contexts in which speaker commitment to the embedded content is not ensured. A paradigm example is when a journalist uses discourse parenthetical reports to report the opinions of two parties who disagree with each other. For instance, in example (8), a toy variant of (13), which is discussed below, the speaker uses one report to express the point of view of the employees and another report to express the point of view of the boss, who directly disagrees with the employees.

- (8)
- a. There was an explosion at the factory.
 - b. The employees said that it was the fault of the boss
 - c. but the boss said that it was the fault of the employees.

We cannot infer speaker commitment to the embedded clause of either report by looking at (8) alone, but both reports are nevertheless used discourse parenthetically to offer possible explanations of the explosion. This is a problem for Hunter et al.'s account because the discourse parenthetical use of the reports calls for annotating the reports with Source, but the fact that the embedded content of the reports is not entailed precludes an annotation with Source. A further problem with Hunter et al.'s account is that even in cases such as (6b) and (7b), in which author commitment to the embedded content of a discourse parenthetical report can be inferred, commitment is almost always inferred by using world knowledge to reason about the reliability of the source(s) cited in the attribution predicate of the report. This kind of world-knowledge based reasoning, central to de Marneffe et al. (2012)'s study of veridicality, is generally independent of the reasoning used to determine rhetorical structure. In other words, Hunter et al.'s Source/Attribution distinction does not reflect a *rhetorical* distinction, and thus it fails to use the notion of discourse function from SDRT (or any rhetorical theory) to model the intuitive discourse function of discourse parenthetical reports.

3.2 No commitment

Veridicality entails that the embedded clause of (1b) cannot be related to (1a) via Explanation, at least if there is any doubt about the speaker's commitment to the truth of this clause's content. This is intuitively correct: if the speaker is not entirely confident about the claim that John was out of town, she cannot be confident about the implicature¹⁴ that John's being out of town explains why he wasn't at the party. Still, in uttering (1b) she makes salient the possibility that John didn't come to the party because he was out of town and performs something *like* an Explanation—a hedged Explanation, if you will. This intuition plays an important role in motivating attachment-based treatments of discourse parenthetical reports.

However, some reports that seem in other ways to be discourse parenthetical carry no requirement of speaker commitment to the embedded clause and are not intuitively used to offer the content of the embedded clause as even a potential explanation or answer, etc. This happens when a speaker explicitly denies the content of the embedded clause or otherwise reveals that she thinks it is false. Simons (2007) discusses some such examples using question/answer pairs.

- (9)
- a. Which course did Louise fail?
 - b. Henry, falsely, thinks that she failed calculus. (Simons, example (18))

These examples are complicated. At first glance, the report (9b) seems discourse parenthetical: the content of the embedded clause appears to be what makes the report relevant to the incoming discourse because it is this content that potentially provides an answer to (9a). On the other hand, in (9b), the speaker cannot be taken as actually offering the proposition *Louise failed calculus* as even a potential answer to the question posed in (9a); in fact, she makes it clear that she is committed to

14. A discourse relation that is not explicitly marked but inferred on the contents of its arguments is cancellable, although many instances will be very difficult to cancel.

that proposition's *not* being an answer. As a result, relating the clause embedded under *thinks* to (9a) via the relation Question-Answer Pair would seem inappropriate.¹⁵

One could also use a negated speech or attitude verb (*didn't say, doesn't think*) or a negative verb (*doubts, is skeptical*) to block the associated speech act. Even more interestingly, from the perspective of rhetorical structure, is that we can undermine an intuitively discourse parenthetical reading by stringing multiple discourse units together to form multi-part responses. For instance, a speaker can disengage herself from the embedded content of a report with a Comment, e.g. (10), or a Contrast, e.g. (11) and (12):

(10) Henry said she failed calculus. He always gets things wrong!

(11) Henry said she failed calculus, but he's wrong.

(12) Henry said she failed calculus, but Anna disagrees.

In fact, it can take many discourse turns to determine a speaker's commitment to the embedded content of a report. The excerpt below is from an article in the *New York Times*.¹⁶ The first two paragraphs of the article discuss an earthquake that hit Prague, Oklahoma in 2011 and the destruction that it caused. The excerpt provides the third and fourth paragraphs:

At a packed town hall meeting days later, Ms. Cooper said, state officials called the shocks, including a 5.7 tremor that was Oklahoma's largest ever, "an act of nature, and it was nobody's fault."

Many scientists disagree. They say those quakes, and thousands of others before and since, are mainly the work of humans, caused by wells used to bury vast amounts of wastewater from oil and gas exploration deep in the earth near fault zones. And they warn that continuing to entomb such huge quantities risks more dangerous tremors — if not here, then elsewhere in the state's sprawling well fields.

These paragraphs work together to present possible explanations for the earthquake introduced in paragraphs 1 and 2. Let's simplify the example for the sake of our discussion.

- (13)
- a. In November 2011, a 5.0 magnitude earthquake shook Prague, Oklahoma.
 - b. State officials said that the earthquake was an act of nature that was nobody's fault.
 - c. However, many scientists argued that the quake was caused by wells used to bury vast amounts of wastewater from oil and gas exploration deep in the earth near fault zones.

15. In (9), we can infer a *negative* answer to (9a), namely that Louise did not fail calculus (Groenendijk and Stokhof, 1984). However, I would like to distinguish (9b) from a report such as (9b'): "Henry thinks/said that she didn't fail calculus". In (9b'), the speaker is really offering the embedded content as a negative answer to (9a), albeit a tentative negative answer. In (9b), the speaker has an answer that is completely independent of what Henry thinks or says—that's why she can judge that Henry is wrong. The point of bringing Henry into it, then, is not so much to use him as a source or as evidence for a potential answer, but to indicate that the speaker is aware that Henry might answer the question differently and thinks that Henry is confused. Thus, Henry's commitments are rhetorically central in (9) in a way that they are not in (9b').

16. 'As Quakes Rattle Oklahoma, Fingers Point to Oil and Gas Industry', by R. A. Oppel Jr. and M. Wines. *The New York Times*, April 3, 2015.

(13b) and (13c) oppose two viewpoints, much like (12). At this point in the article, we can only take the author to be presenting two possible explanations.¹⁷ The rest of the (lengthy) article makes it clear that the authors side with the scientists, however. The following paragraph, taken from the same article, reveals this endorsement:

But in a state where oil and gas are economic pillars, elected leaders have been slow to address **the problem**. And while regulators have taken some protective measures, they lack the money, work force and legal authority to fully address **the threats**.

At this point, the discussion shifts from trying to find an explanation for the earthquake to a discussion of why the local government has been so slow to address the problem brought up by the scientists. The assertions are no longer in the scope of reports, so we are dealing here with the authors' commitments. The use of definites such as *the problem* and *the threats* presuppose the existence of the problem/threat that the scientists introduce in (13c). The larger discussion of why the local government has been so slow in addressing the problem also presupposes that the question of what the problem is has been answered. As this paragraph reflects the authors' commitments, we can infer that they have sided with the scientists, and thus have no commitment to the embedded content of (13b).

We can make a similar point with (1). Imagine an utterance of (1) followed by either (14) or (15) (but not both):

- (14) She must be covering for him, though, because I spotted them together at the market this morning.
- (15) But he was only an hour away. He could have come if he'd wanted to. There must be some other reason.

In judging (1) alone, an interpreter would be entitled to infer that the speaker is committed to the possibility of John's having been out of town and is using this possibility to provide a possible explanation of why John didn't come to her party. However, once the speaker continues with either (14), which denies that the potential explanandum holds, or (15), which denies that the potential explanandum actually explains John's absence, an interpreter is no longer justified in inferring even a hedged Explanation relation between (1a) and (1b). The attitudes that the speaker takes towards the truth of the embedded content of a report and its discourse function are not revealed by considering only the discourse move immediately preceding the report.

Starting definitions of discourse parenthetical reports, like that offered by Simons (2007) or the attachment-based proposals offered by various rhetorical frameworks, are driven by the intuition that the embedded contents of these reports play a certain discourse function. Intuitions are primed by considering pairs consisting of a report and a single, preceding discourse unit. The problem brought out by examples like (10)-(13) is that if we place these pairs in a larger discourse context, our intuitions about the discourse function of the report, and what inferences are licensed by it,

17. The fact that scientists are often taken to be more reliable sources than state officials with regard to causes of natural disasters might lead a reader of (13) to suspect that the writers side with the scientists and wish to adopt their explanation. However, we do not get that from looking at the discourse structure of (13) alone. This kind of world knowledge, discussed by de Marneffe et al. (2012), also plays an important role in the interpretation of reports in discourse, but the focus of this article is on how the rhetorical structure of discourse, and the anaphoric connections between discourse units, affects interpretation.

can change significantly. But then what should we say about the reports in such cases? Are they discourse parenthetical or not? What features shall we use to decide? Attachment-based accounts provide no clear answer.

Simons' discussion waffles on these questions as well. Of examples like (9) (her (18)), she says that: "main point content cannot be identified with the content of either the subordinate clause alone or the main clause. Rather, main point content emerges from the interaction between the subordinate clause content and the attitudes to that content expressed by the other predicates used." This undermines her definition of discourse parenthetical reports. In the beginning, she defines a discourse parenthetical report as one whose embedded clause conveys the main point of the utterance, but later, she seems to suggest that reports like (9b) should count as discourse parenthetical. Although she notes this tension, she does not go on to work it out, as her focus is on other issues, so the ensuing discussion does not point to a path down which rhetorical accounts of discourse parentheticals can go.

3.3 Distinct functions

In the previous sub-section, we discussed examples in which a report is offered as a response to a single discourse unit (e.g. *Which course did Louise fail?*), but in which the full discourse contribution of the report—the information that determines what conclusions can be drawn from the report—can only be understood by looking at chunks of discourse involving multiple discourse units. In this section, we note that the other direction is possible as well: sometimes a single report is a coherent response to more than one discourse unit. In other words, the embedded clause and the attribution predicate might each have a discourse function relative to the incoming discourse, but these functions are distinct and we might not catch both of them if we only consider a single incoming discourse unit (such as a question).

Let's return to (13), but alter it slightly to bring it more in line with the original excerpt.

- (16)
- a. In November 2011, a 5.0 magnitude earthquake shook Prague, Oklahoma.
 - b. State officials said that the earthquake was an act of nature that was nobody's fault.
 - c. Many scientists disagree.
 - d. They argue that the quake was caused by wells used to bury vast amounts of wastewater from oil and gas exploration deep in the earth near fault zones.

(16c) describes the scientists' attitudes, and (16d) intuitively elaborates on this description. Yet (16d) can only be an Elaboration on (16c) if we treat the attribution predicate as rhetorically relevant; the embedded clause alone does not talk about the scientists' attitudes, but only about the object of their attitudes. Thus locally—that is, looking at the report in (16d) as a response or follow up to (16c) alone—the main point of the report seems to be to provide more information about what the scientists hold and the report is not obviously construed as discourse parenthetical. However, as we saw in the last subsection, if we look at the way (16b-d) (or the associated part of the original excerpt) is functioning in the larger discourse, we see that the embedded clause of (16d) serves a higher discourse function, namely that of conveying a possible explanation of why the earthquake occurred. On this level, it is the embedded content that seems more rhetorically important or that carries the main point.

Once again, we see that the effect of reports on discourse structure can be subtle and diffuse and the intuitions about main point and discourse relevance triggered by looking at pairs of utterances, e.g. (1), do not carry us through to a full account of reports in discourse.

4. Modeling Reports in Discourse

The discussion from §3 highlights three issues that a theory of discourse parenthetical reports needs to address. First, because of *hedged commitments*, the use of a discourse parenthetical report can be taken to be at most a speech act that conveys the *possibility* of a rhetorical connection between two eventualities. Second, because of examples of *no commitment*, even an act of hedging seems too strong in many cases—a speaker might use a report that at first seems discourse parenthetical only to later reveal a lack of commitment to the embedded content. Third, because of *distinct functions*, both parts of a discourse parenthetical report can enter into discourse relations with the larger context; in other words, both can serve a discourse function in the sense defined by a theory of rhetorical structure. As these issues remain unresolved, our working definition of a discourse parenthetical report remains vague: a discourse parenthetical report is one in which the embedded clause plays a major role in making that report relevant to the preceding discourse.

The aim of this section is to clarify what it is for a report to be discourse parenthetical and to develop a model of the discourse contributions of such reports. The model that I favor is presented in §4.2. But first, §4.1 lays out my reasons for rejecting a very different, but seemingly attractive, alternative model.

4.1 Separate layers of information?

The problems of *hedged commitments* and *no commitment* stem largely from a conflict between the intuitive discourse function of a discourse parenthetical report and the kind of speaker commitments that such a function normally entails. An appealing hypothesis, then, is that discourse structures should track information about discourse relations and information about a speaker's attitude to those relations in two separate dimensions (cf. Potts (2005)'s multi-dimensional account of (not-) at-issue content). The PDTB annotation approach suggests, although it does not entail, one such model. Recall that the content of the attribution predicate of a discourse parenthetical report never contributes to the argument of a discourse connective in the PDTB. Nevertheless, this information is stored alongside the annotations and could in theory be exploited during the interpretation of the annotations. Such a two-dimensional approach might also work to model Simon's claim that the two clauses of a discourse parenthetical report "convey two different types of content" (Simons 2007, p. 1053), although because Simons does not provide a model or a formal notion of discourse function, she is not committed to such an approach.

The problem with a two-dimensional model that treats the attribution predicate and the embedded clause as conveying two different types of content is that it does not square with the observation of *distinct functions*. If both parts of a discourse parenthetical report can enter into discourse relations, i.e., both parts make the same type of contribution to discourse structure, then whatever differences may exist between the types of content conveyed by the two clauses, it is not a difference that motivates a model in which the attribution predicate contributes the whole of its content to one

dimension or layer of discourse structure while the embedded clause contributes the whole of its content to another.¹⁸

An alternative approach would be to posit two layers of information, one that tracks rhetorical relations and one that tracks commitment to those relations, but allow the attribution predicate to contribute to both layers. Danlos and Rambow (2011) proposes an account along these lines. The general idea is to annotate the relations that intuitively hold between discourse units without regard for the speaker’s commitment to the content of those units. Then, as each labelled attachment is constructed, formulas are added to the annotation that contain information about the speaker’s attitude towards each argument of that relation. These formulas take the form: $f(e, s) = Mod, Pol$, where s is an agent of a propositional attitude, e is the eventuality that serves as the object of the attitude, and f is a propositional attitude function that maps a pair (e, s) to a judgment about e . Following Saurí and Pustejovsky (2009), these judgments have a modal component (*Mod*)—which can take the value certain (CT), probable (PR), possible (PS), or unknown (U)—and a polarity component (*Pol*)—positive (+), negative (-), or unknown (u). Finally, each discourse relation in the annotation structure is subscripted with a source. In general, the source will be the writer or speaker, but in the case of reports, the source can be the source of the reported content.

To illustrate, in (17), the Narration relation that intuitively relates Bill and Jane’s attending dinner and their going dancing would be attributed to Bill.

- (17) [Jane had dinner at my place on Thursday,] _{α} [and Bill said that] _{β_{att}} [she went dancing afterwards.] _{β}

This would (skipping a few details that are irrelevant here) yield the following, two-part annotation:

- (i) $\text{Attribution}_{Wr}(\beta_{att}, \beta), \text{Narration}_{Bill}(\alpha, \beta)$
(ii) $f(e_\alpha, Wr) = CT+ \wedge f(e_\beta, Wr) = Uu \wedge f(e_\beta, Bill) = CT+$

‘ $\text{Attribution}_{Wr}(\beta_{att}, \beta)$ ’ means that the writer is committed to an Attribution relation between β_{att} and β and ‘ $\text{Narration}_{Bill}(\alpha, \beta)$ ’ means that Bill is committed to a Narration relation between α and β . ‘ $f(e_\alpha, Wr) = CT+$ ’ means that the writer is certain that e_α holds (similarly for ‘ $f(e_\beta, Bill) = CT+$ ’), and ‘ $f(e_\beta, Wr) = Uu$ ’ means that both the truth of e_β and the author’s attitude towards its truth are unknown.¹⁹

While this proposed annotation captures many intuitive features of (17), the proposed account fails to make the right predictions. First, (17) supports neither the inference that Bill is committed to $\text{Narration}(\alpha, \beta)$ nor the inference that he is committed to e_β having occurred; it justifies at most the conclusion that the speaker is committed to Bill’s being committed to these things—the speaker could be making all of this up just to get Jane in trouble. Second, there is a question of why we

18. Maier and Bary (2015), following work by Potts (2005), develops a different kind of two-dimensional model in which the embedded clause of a discourse parenthetical report contributes to one dimension and the entire content of the report (including the embedded clause again) contributes to another. Maier and Bary do not explore a discourse-based account of parenthetical readings, opting rather for an ambiguous semantics for reports that distinguishes between parenthetical (two-dimensional) and non-parenthetical (one-dimensional) reports. For this reason, I will not explore their account in the current paper; nevertheless, I suspect that the troubles brought out by *distinct functions* will apply to their account as well, as such examples count against a simple binary distinction in which only the matrix clause or only the embedded clause is active in a given interpretation.

19. It is unclear whether the speaker’s attitude towards the attribution predicate is recorded. It does not appear to be, though it is not obvious in Danlos and Rambow (2011) why it would be excluded, since it figures in discourse relations.

would want to track Bill's commitments in the first place. The units β_{att} and β are only relevant to α insofar as the speaker is committed to there being a relation between α and some part of the report. Discourse structures track speech acts of the *speaker* because it is the speaker who chooses the discourse.

It is also unclear why annotations of discourse structure must track the distinction between a speaker's thinking an eventuality is possible and her thinking that it is probable. Distinctions like this will in general be entailed by the semantics of the embedding verb chosen (*know* is factive, etc.) or by the way that the report is used in the larger discourse context. This information doesn't require an extra layer of annotation on discourse structure. Plus, even if we wanted to make such modal distinctions, they could be captured by appending modal operators to the relations (e.g., $\diamond\text{Explanation}(\alpha, \beta)$). We need modal relations in any case to handle responses like: *Well, he was out of town, so maybe that's why he didn't come*, in which the speaker expresses her commitment to the potential explanans but hedges her commitment to the relation itself. Tracking commitments with modalized relations also entails the requisite commitments to the arguments themselves: if a speaker is committed to $\diamond R(\alpha, \beta)$ for a veridical relation R , then veridicality entails that her level of commitment to β be at least $\diamond\beta$, which is consistent with either full commitment to β or commitment only to $\diamond\beta$. Thus, by putting the modal on the relation, we get the right results for both the relation and the arguments. By contrast, merely putting the modal on the argument itself (which is equivalent to what Danlos and Rambow propose) does not achieve what is intuitively required—a modal *relation*. The import of the report in (1) is not correctly captured as $\text{Explanation}(1a, \diamond\textit{he was out of town})$, but as $\diamond\text{Explanation}(1a, \textit{he was out of town})$.

The more general point is that if a speaker uses a discourse parenthetical report to hedge her commitment to a discourse relation $R(\alpha, \beta)$, then, as explained in §3.1 and §3.2, she is not performing a speech act whose content is $R(\alpha, \beta)$. Annotating (1) with the formula $\text{Explanation}(1a, \textit{he was out of town})$ and then separately adding the information that the speaker is not fully committed to this content simply isn't the right thing to do in a rhetorical theory, because there is no speech act of *Explanation* that has taken place; the speech act just is a proposal of a possible explanation. One has to take the semantics of speech acts into account in calculating what speech act took place; consequently, the veridical nature of veridical relations and their discourse function cannot actually be separated as a two-dimensional or two-layer theory would propose.

The natural thing to do within a rhetorical theory is to see the way that speakers use discourse parenthetical reports as speech acts in their own right by assigning them a different kind of discourse relation and then associating these relations with their own semantics. I develop such an account in the next subsection.


4.2 An integrated model

I develop my account of reports in discourse by extending SDRT and modifying attachment-based analyses of discourse parenthetical reports to handle the three issues outlined in §3. Let's begin with the problem of *distinct functions*, that is, the fact that both parts of a discourse parenthetical report r can be rhetorically relevant by simultaneously being linked to distinct discourse units in the discourse graph for the discourse preceding r . To capture these data, we must allow both parts of a report to contribute to arguments of discourse relations, which in turn requires that we distinguish the attribution predicates of all reports from their embedded clauses by segmenting the two parts as separate discourse units. Thus a given report r will contribute two segments to discourse: the

attribution predicate, whose content entails $\exists p.x \text{ said } (thinks...) \text{ that } p$ for some agent x supplied by the report, and the embedded clause, which specifies the content of p .

Segmenting reports this way is somewhat controversial, because the attribution predicate does not denote a precise eventuality on its own; it merely tells us that *something* was said (or thought, etc.). Only when combined with the embedded clause does it adequately specify the requisite speech event (or attitude). This gives some merit to the PDTB and CDT choice to not treat reports as contributing two distinct segments despite the fact that reports do convey two eventualities: that of the saying or thinking, etc. and that described by what was said or thought, etc. The benefits of segmentation outweigh the complications, however. We have too much evidence that both clauses of a discourse parenthetical report can be discursively relevant to ignore one or the other. Plus, the idea of creating separate discourse units for intimately intertwined contents is already required for treating other phenomena such as presupposition in SDRT. In the case of presupposition, there is no syntactic unit that circumscribes a discourse unit for the presupposition; nevertheless, a separate unit for the presupposition is semantically motivated and required in the discourse structure.²⁰

The next step in modelling the discourse function of discourse parenthetical reports is to determine how the attribution predicate and embedded clause are related to one another. Recall that Hunter et al. propose a special relation, Source, to handle discourse parenthetical readings. In contrast, I propose that all reports be annotated with Attribution, regardless of whether they receive a parenthetical reading or not. That is, given two segments, α and β , where α labels the attribution predicate of a report r and β labels the embedded clause, I propose that r contributes the following structure to a discourse graph:



Structurally, Attribution is a subordinating relation, which reflects the syntactic structure of discourse parenthetical reports.

For its semantics, suppose we have an instance of $\textit{Attribution}(\alpha, \beta)$ for some discourse units α and β , and let K_α be the DRS that represents the content of α , and K_β , the DRS that represents the content of β (see Kamp and Reyle (1993)). K_α should entail that some agent x stands in some attitude A to a proposition p . For example, the attribution predicate of *Jill said John was out of town* entails the content $\textit{said}(Jill, p)$ for some proposition p (i.e. *Jill said something*). The formula $\textit{Attribution}(\alpha, \beta)$ is then true just in case the content of K_α is true and the content of K_β specifies the content of the proposition p . Or in dynamic semantic terms, update with $\textit{Attribution}(\alpha, \beta)$ at a world w and assignment f extends f to an assignment g that classically satisfies $\textit{Attribution}(\alpha, \beta)$ at w just in case w and f likewise support update with K_α and the resulting assignment g assigns p a subset of the w -accessible worlds w' that support update with K_β given the assignment g . More formally:

- $(w, f) \llbracket \textit{Attribution}(\alpha, \beta) \rrbracket (w, g)$ iff $g \supset f$ and $\exists x \exists p. K_\alpha \models A_{K_\alpha}(x, p)$ and $(w, f) \llbracket K_\alpha \rrbracket (w, g)$, and $g(p) \subseteq \{(w', k) : (w', g) \llbracket K_\beta \rrbracket (w', k)\}$

20. Thanks to Laure Vieu for discussions on this point.

where A_{K_α} is an attitude predicate determined by the embedding verb, e.g. *said*, *thought*, etc., contributed by α .

To capture the discourse function of discourse parenthetical reports, we need to model the relation between the report and the incoming discourse context. I propose that this relation has three features, which come together to make a report discourse parenthetical. First, the embedded clause must be related directly to a discourse unit introduced in the discourse preceding the report. This feature is common to all attachment-based solutions. Second, the relation must be distinct from the relation via which the attribution predicate is related to the incoming context (if the attribution predicate is so related). Third, the relation connecting the embedded clause to the preceding context will be a modal relation indicating a *hedged commitment* from the speaker.²¹ The function of (1b), for example, is to provide a *possible explanation* of John's absence from the party.

Modal relations are triggered by reports as follows. Suppose we have an instance of Attribution with the form $Attribution(\alpha, \beta)$ and suppose that there is a discourse link whose second argument is β and whose first argument, γ , is found in the discourse preceding α . The link between γ and β will be labelled $\diamond R$ for some discourse relation R . More precisely, we can add the rule PA, for *parenthetical Attributions*, to the logic of SDRT:

$$\text{PA: } (\exists e. e(\alpha, \beta) \wedge l(e) = \textit{Attribution}) \rightarrow (\exists \gamma \exists e' (\gamma <_t \alpha \wedge e'(\gamma, \beta)) \rightarrow \exists R (l(e') = \diamond R))$$

In words, PA says that if there is an edge e (in the discourse graph) that connects the nodes α and β and e is labelled *Attribution*, then if there is another discourse unit γ that is introduced in the discourse prior to α and if there is an edge e' that connects γ and β , then e' will be labelled $\diamond R$ for some discourse relation R .

Visually we can represent the import of PA as in Figure 1. Suppose that we have $Attribution(\alpha, \beta)$ for some units α and β , and that we want to attach α or β or both to the preceding discourse. Suppose further that γ is an accessible discourse unit in the structure for the discourse preceding α . Then, a connection from γ to α will proceed as normal in SDRT and will be labelled with some relation R . However, a connection from γ to β will license a label $\diamond R'$ for some R' .²²




Figure 1: Relations to units inside Attributions

-
21. Certain verbs, notably factives, can strengthen this contribution to discourse structure, and see §4.2.5 for a discussion of when discourse structure can itself strengthen the connection.
22. R' might be either subordinating or coordinating. In SDRT, subordinating relations are usually represented with down arrows, \downarrow , while coordinating relations are represented with right arrows, \rightarrow . Figure 1 is meant to be neutral with regard to the subordinating/coordinating distinction for R' , though later in the paper, when a particular relation is specified, diagonal edges like that in Figure 1 are used for subordinating relations when a pure down arrow would overlap another down arrow in the graph.

The interpretation of a formula $\diamond R(\alpha, \beta)$ for a veridical relation R will be derived from the interpretation of R . Veridicality entails that a formula of the form $R(\alpha, \beta)$ (for veridical R) can be true (i.e. dynamic update yields a non-empty set of world-assignment pairs that classically support $R(\alpha, \beta)$) if and only if (i) the contents of α and β , which are represented in the DRSs K_α and K_β , respectively, are true, and (ii) the set of constraints determined by R are satisfied. Explanation, for example, imposes the constraints that the eventuality described by β begin before the eventuality described in α and that the former is the cause of the latter in the world of evaluation, and so on. More formally, the satisfaction schema for a veridical relation R is as follows (see page 156 of Asher and Lascarides (2003)):

- $(w, f) \llbracket R(\alpha, \beta) \rrbracket_M(w', g)$ iff $(w, f) \llbracket K_\alpha \wedge K_\beta \wedge \phi_{R(\alpha, \beta)} \rrbracket_M(w', g)$

where \wedge is dynamic conjunction, M is a model, and $\phi_{R(\alpha, \beta)}$ the set of constraints determined by the particular relation R . A formula of the form $\diamond R(\alpha, \beta)$ for veridical R will thus be true (i.e. dynamic update yields a non-empty output) under the following conditions:

- $(w, f) \llbracket \diamond R(\alpha, \beta) \rrbracket_M(w', g)$ iff $\exists w'' . w \mathcal{R} w''$ such that $(w'', f) \llbracket R(\alpha, \beta) \rrbracket_M(w', g)$

That is, dynamic update with $\diamond R(\alpha, \beta)$ at a pair (w, f) will yield a non-empty output that classically supports $\diamond R(\alpha, \beta)$ just in case there is a world w'' that is epistemically accessible to w that together with f supports update with $R(\alpha, \beta)$.

A consequence of this semantics is that from an instance of a veridical relation of the form $\diamond R(\alpha, \beta)$ we can conclude at most that the speaker is committed to $\diamond\alpha$ and $\diamond\beta$ unless α or β are asserted elsewhere in the discourse. This resolves the problem of *hedged commitments*. At the same time, we can infer that the speaker is committed to *at least* $\diamond\alpha$ and $\diamond\beta$; $\diamond R(\alpha, \beta)$ cannot hold if either $\diamond\alpha$ or $\diamond\beta$ fail to hold. A speaker cannot (consistently) commit to the possibility that a discourse unit answers a question, provides an explanation, continues a narrative, and so on, if she is committed to its content being false.

A further consequence of the solution I have proposed is that if a speaker is committed to Attribution(α, β) and $\neg\beta$, then the report that contributes Attribution(α, β) to the logical form of the discourse is *not* discourse parenthetical in my account. Consider (18), an example of *no commitment* in the sense of §3.2:

- (18) John didn't come to my party. Jill said he was out of town, but she's wrong. He's just avoiding me.

It's tempting to say that the report is discourse parenthetical in this example just as it is in (1). After all, the discourse properly includes the content of (1). However, there is no sense in which the embedded clause of the report is serving the function of explaining or even starting to explain why John didn't come to the party; the speaker is not connecting the embedded clause and the preceding discourse but rather explicitly cancelling the implication that there is a positive relation between them. The report feels discourse parenthetical because the embedded clause *might have been* used to perform a discourse function, but the speaker's point is that it *actually* doesn't. Thus examples of *no commitment* are not discourse parenthetical in my account.

I turn now to five examples that illustrate the solution that I have proposed and explain in more detail how they solve the problems of *hedged commitments*, *no commitment*, and *distinct functions*. The five examples represent the most frequent problem cases that I have encountered in annotating or studying well over 100 newspaper articles from *The Wall Street Journal* (in examination of the PDTB and annotation work on the DisCoR corpus) and *The New York Times* (through personal research).

4.2.1 Single report

Single report examples are the sort of examples that are usually used to motivate discourse parenthetical readings. They take the form of a single report, e.g., (1b), that is understood as a response to a preceding discourse unit, e.g., (1a). The speaker uses the report to bring out the possibility that the embedded clause serves the relevant discourse function. Single reports yield graphs like Figure 2.




Figure 2: Single report

Figure 2 reflects the fact that the function of (1b) is to suggest the possibility that John didn't come because he was out of town. It also shows that in examples of single reports, the attribution predicate is rhetorically inert. Note, however, that while rhetorically inert in the sense that it does not attach to a preceding discourse unit, the attribution predicate has a significant impact on the logical form of the discourse because it weakens the inferred discourse relation. Recall that Simons argued that in examples such as (9b), both parts of the report work together to yield the response given to the preceding question (e.g., 9a). In an important sense, however, the attribution predicate and embedded clause *always* work together to yield the full content of the response. The account that I have proposed models this joint contribution by positing that discourse parenthetical reports give rise to modal discourse relations.

4.2.2 Contrasting viewpoints

Newspaper articles also frequently use reports to contrast two opposing viewpoints, where the author takes a stand on neither one. Imagine the following continuation of (1):

- (19) John didn't come to my party. Jill said that he was out of town, but Bill said that he's just avoiding me. (So I don't know.)

The reports in these examples are used as single reports are; it's just that more possibilities are discussed. We therefore get two edges labelled with $\diamond \textit{Explanation}$: one whose endpoint is in the scope of the first Attribution and one whose endpoint is in the scope of the second, as shown in Figure 3.

In addition to these two edges, we also have a Contrast between the two reports. With contrasting viewpoints, the Contrast edge will always link the attribution predicates in cases where the viewpoints are described using two reports in which the attribution predicate takes syntactic scope over the content of the reported speech or attitude. The idea, following Hardt (2013), is that if two embedded contents satisfy the requirements of Contrast, then the belief or commitment state of an agent who is committed to one embedded clause will likewise be in contrast with the belief or commitment state

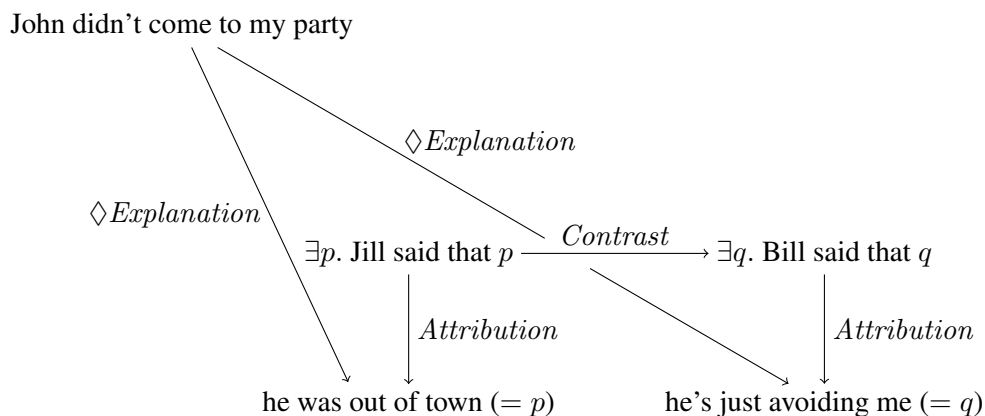


Figure 3: Contrasting Viewpoints

of an agent who is committed the other. Representing the contrast at the level of the agents' attitudes or commitments captures the full contrast.²³

4.2.3 Mixed signals

Mixed signal examples also involve multiple reports, but arise when the attribution predicates of the reports suggest one relation, while the embedded clauses suggest a very different relation. These cases thus involve examples of *distinct functions* in the sense of §3.3. (20) is an example:

- (20) John told his parents that he had an accident (so they felt bad for him). However, the police later said that he had been drinking (so now they're mad!).

The attribution predicate of the second report enters into both a Narration and a Contrast relation with the attribute predicate of the first. Both relations are explicitly marked, by *later* and *however*, respectively, and could be rhetorically significant in the larger discourse. For example, the Narration might be relevant because John's parents felt bad for him at first but later came to be angry at him, as suggested by the clauses in parentheses. As for the embedded clauses, the embedded clause of the second report suggests a Background relation with that of the first report because it suggests that John had been drinking *before* he had the accident. The temporal requirements of $\text{Background}(\alpha, \beta)$ exclude the possibility of $\text{Narration}(\alpha, \beta)$. Thus not only do the two attribution predicates and the two embedded clauses of (20) enter into distinct discourse relations, they enter into *incompatible* discourse relations, making these cases tricky to model in the attachment-based accounts discussed in §2.

Figure 4 illustrates my proposal applied to examples of mixed signals, and shows that these examples are similar to examples of single reports, except that the incoming link to the second embedded clause comes from inside of another report.²⁴ Still, because the link (here Background) crosses

23. See Hardt (2013) for a detailed argument of why the attribution predicate should contribute to the second argument of Contrast in (6). Hardt's argument is slightly different from the one made here, but in the same vein.

24. Figure 4 does not include nodes for the remarks in parentheses in (20) because adding them would require a discussion of the annotation of *complex discourse units* in SDRT. A complex discourse unit is a group of simpler discourse units that work together to form a single argument to a discourse relation. For example, in *Jane is mad at Fred because he*

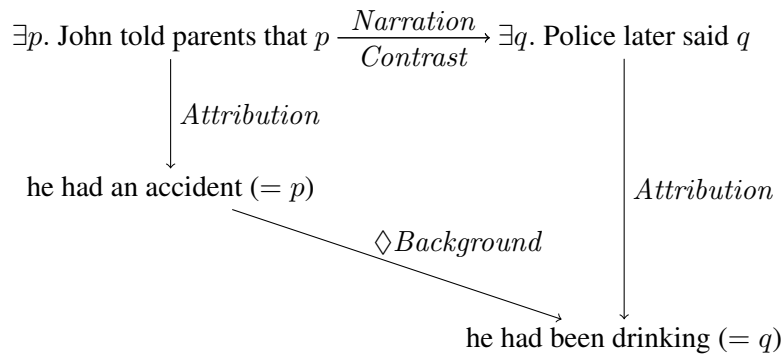


Figure 4: Mixed Signals

the boundary of a report context, we get an inference to a weaker relation, namely \diamond Background, between the embedded clauses.

4.2.4 Speaker disagreement

Examples of speaker disagreement are a subset of the examples of *no commitment* discussed in §3.2 in which the speaker explicitly disavows the embedded content of a report that might otherwise have seemed discourse parenthetical. (11), repeated here, is an example.

- (11) (*Which course did Louise fail?*) Henry said she failed calculus, but he's wrong/but she didn't.

As explained above, I take the stand that reports used in this way are *not* discourse parenthetical precisely because a speaker who uses them does not offer the embedded content as even a possible Answer, Background, Explanation, or Narration, etc. to a preceding discourse unit γ . A speaker might at first appear to be using a report this way—in the sense that were we to look only at the report and the requisite preceding unit γ , we would understand the report parenthetically—but as soon as she explicitly rejects the embedded content, she makes it clear that she is not.

I propose that we analyze examples of speaker disagreement as involving the cancellation of a defeasible inference. In SDRT, when a discourse relation is not explicitly marked, the inference of the relation is defeasible. This is why SDRT assigns discourse relations non-monotonic or probabilistic semantics and stresses the need to consider the larger discourse context in making the ultimate choices about how a pair of discourse units is related. Sometimes, a speaker can explicitly cancel the inference of a discourse relation that she, given the knowledge that the interlocutor can be assumed to have, might expect her interlocutor to draw:

- (21) A man died today after being run over by a fire truck. However, it wasn't the fire truck that killed him.

broke her sunglasses and lost her favorite book, it is intuitively both *he broke her sunglasses* and *lost her favorite book* that together explain why Jane is mad at Fred. In an annotation of (20) that included the remarks in parentheses, all of the units of the first sentence would be grouped together and together they would stand in a Contrast/Narration relation to the group of units formed by all of the units in the second sentence. Further discussion of complex discourse units would introduce complications that would take us too far afield here.

- (22) You need to take a bath, do your homework and pack your backpack, though not necessarily in that order.

In (21), the causal inference—which is supported by, among other things, world knowledge concerning the relative size and weight of fire trucks—is cancelled leaving only the temporal interpretation of *after*. In (22), the inference that events take place in the order presented is cancelled. Along these lines, I propose that in examples like (11), a speaker’s explicit disavowal of the embedded content of a report that might otherwise be understood as discourse parenthetical amounts to a disavowal of the relevant relation as well.

If reports like that in (11) are not discourse parenthetical, then what is their relevance to the discourse? I suggest that speakers use them to provide background information relevant to the general discussion. Let r be a report that at first appears to be discourse parenthetical and thus to provide an answer, explanation or so on, to another discourse unit γ . Suppose as well that the speaker explicitly cancels the implicated relation between the embedded clause of r and γ . The point of including the report despite ultimately rejecting its embedded content is to indicate that the speaker is not discussing γ in ignorance of certain information possibly pertinent to γ . If γ is a question, for example, the speaker indicates that the question is still an open one (at least at time that the speaker provides her own answer, if she provides one)—in (11), the question has not been answered by anything Henry said. Similarly for other relations: in the example from Figure 5 below, the point is to show that what Jill said did not in fact explain John’s absence. So the report is really Background information on the speaker’s utterance of γ and on what is going to come. I therefore propose that such reports be annotated as shown in Figure 5.

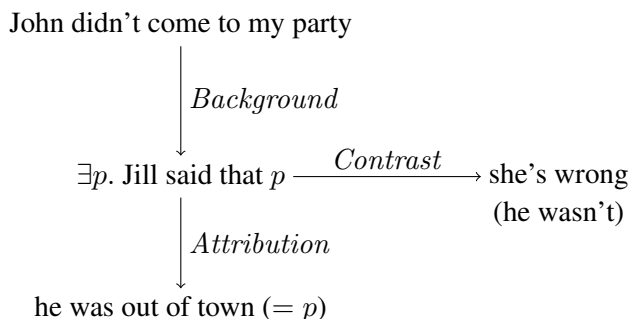


Figure 5: Speaker Disagreement

Note that whether the contrast is with *she's wrong* or *he wasn't*, the contrast attaches to the attribution predicate *Jill said that*. The content *he wasn't* cannot attach to *he was out of town* because doing so would entail the conjunction *he was out of town and he was not out of town*, which is inconsistent. Even if we tried to posit a \diamond Contrast relation to the embedded clause, we would get *possibly, he was out of town and he was not out of town*, which is no better. Attaching *he wasn't* to the attribution predicate also correctly reflects the fact that in saying *he wasn't*, the speaker is saying that her commitments are in conflict with Jill’s commitments. In this case, therefore, the attribution predicate, which tells us that we are talking about Jill’s commitments, is rhetorically relevant to the contrast.

4.2.5 Strengthening

Examples of strengthening are examples in which it becomes clear from the discourse context that a speaker who uses a discourse parenthetical report is committed to the content of the embedded clause. Suppose for the sake of illustration that the authors of (16) had never considered the point of view of the state officials, so that only the point of view of the scientists had been considered as in (23).

- (23)
- a. In November 2011, a 5.0 magnitude earthquake shook Prague, Oklahoma.
 - b. Scientists said that the quake was caused by wells used to bury vast amounts of wastewater from oil and gas exploration deep in the earth near fault zones.

This would yield a graph in which the embedded clause of (23b) is attached to (23a) with \diamond Explanation. Now recall from the discussion in §3.2 that the authors ultimately endorse the point of view of the scientists. This (ignoring the point of view of the state officials) is a pure case of strengthening. In such examples, no revision of the \diamond Explanation relation is required: it is semantically compatible with its arguments, including the embedded clause of (23b), being true. Moreover, the truth of the embedded clause will be entailed by later veridical connections to the report. Consider (24a), which comes after the full version of (23b) in the original article (see footnote 16) and (24b-e), which come a full paragraph after (24a).

- (24)
- a. And they warn that continuing to entomb such huge quantities risks more dangerous tremors.
 - b. But in a state where oil and gas are economic pillars,
 - c. elected leaders have been slow to address the problem.
 - d. And while regulators have taken some protective measures,
 - e. they lack the money, work force and legal authority to fully address the threats.

Combining (23) and (24), we would get Continuation(23b,24a) at the level of the attribution predicates, and Contrast(24a,[24c,24e]), where [24c,24e] is a *complex discourse unit* whose constituents, which are related to each other via Continuation, together form the second argument to the Contrast. ((24b) and (24d) would be related to their associated arguments via Background, but that's not central here.) Because (24c) and (24e) enter into a Contrast relation with the attribution predicate of (24a) and because they are not in the scope of an Attribution context, veridicality entails that they must both be true. This in turn entails that the definites *the problem* and *the threats* have antecedents, which, finally, requires that the scientists' judgments be true.

Returning now to (16), repeated here for convenience, the strengthening of the scientists' report entails in addition a rejection of the possible explanation between what state officials said and the earthquake.

- (16)
- a. In November 2011, a 5.0 magnitude earthquake shook Prague, Oklahoma.
 - b. State officials said that the earthquake was an act of nature that was nobody's fault.
 - c. Many scientists disagree.
 - d. They argue that the quake was the result of continued oil and gas exploration deep in the earth near fault zones.

This example is an example of strengthening plus *no commitment* in the sense of §3.2. While the disagreement is not made explicit in the way that it is in (11), handling this disagreement will involve a revision procedure for the attachment of (16b) along the lines of that outlined for speaker disagreement (see §4.2.4).²⁵

4.3 The Right Frontier

The solution offered above violates the Right Frontier Constraint as defined in SDRT. The Right Frontier of a discourse graph or tree is a set of nodes in the graph that is, as its name suggests, found along the rightmost side of a discourse graph (where arcs extend down from the root or to the right of other nodes in the graph). As a discourse evolves, the set of nodes on this frontier will change. The point of the Right Frontier is to track the set of salient, accessible nodes in the graph at a given point; the Right Frontier *Constraint* is a constraint that a new discourse move m be attached to one of the nodes that is on the Right Frontier at the time that m is added to the discourse.²⁶

The Right Frontier for a new discourse move m always includes the discourse unit that immediately precedes m , i.e. move $m - 1$. This move is called *Last* because it is the textually last move that was added to the discourse graph to which m needs to be attached. Aside from *Last*, the Right Frontier includes any node that is connected to *Last* via a chain of subordinating relations, although details on subordinating versus coordinating relations are not important here because the violation engendered by discourse parenthetical reports concerns *Last*. In our analysis of (1), let χ be the segment (1a), α the attribution predicate in (1b), and β , the embedded clause of (1b). Then χ textually precedes α , and α textually precedes β . When it is time to attach β to the discourse graph for χ and α , α is *Last*; χ is neither *Last* nor attached to *Last* via a subordinating relation (or any relation) and so χ is not on the Right Frontier for β . Nevertheless, β attaches to χ , giving rise to a violation.

Rather than adjust the solution that I have offered in an effort to preserve the Right Frontier Constraint, I think the correct reaction in this case is to acknowledge that discourse parenthetical reports engender a genuine, though restricted, violation of this constraint as traditionally conceived.²⁷ In annotation projects in SDRT, we have seen consistently that violations of the Right Frontier Constraint are almost inevitably related to speech and attitude reports and arise because annotators want to make a direct attachment from the preceding discourse context into the embedded clause of a report. This is not so surprising after all: the Right Frontier Constraint was developed as a constraint on how a speaker presents her own commitments in monologue. When a speaker chooses to use someone else's commitments to accomplish this task, she must *first* set up the other speaker's commitments and *then* use them.

I propose an adjustment to the Right Frontier Constraint that takes this two part procedure at face-value by first setting up sub-graphs for Attributions and then letting either argument to the Attribution

25. The fact that the disagreement is made via an explicit move in speaker disagreement cases makes those cases slightly more complicated than I presented them in §4.2.4. Explicit rejection involves an explicit discourse move of *Correction*—a move that is absent in examples like (16). Going through all of the steps to handle an example like (11) would therefore require going through the semantics of *Correction*, a discussion that would take us too far afield. See Venant et al. (2014) for a pertinent discussion.

26. See Asher (1993) for a discussion of Right Frontier violations (which are rare and marked).

27. Venant (2014) develops a proposal to save the Right Frontier Constraint in examples involving discourse parenthetical reports. I do not think this is necessary, or even desirable, for reasons discussed in this section. However, Venant's discussion and proposal are interesting, and he tackles the topic of multiply-embedded Attributions, which I do not broach here.

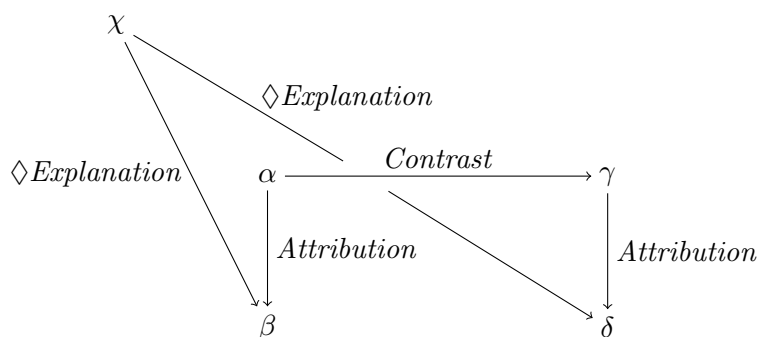


Figure 6: Contrasting Viewpoints (structure only)

attach to the Right Frontier of the incoming discourse graph. As an illustration, consider the graph for Figure 3. I have repeated the graph in Figure 6, but have replaced each node with a discourse unit variable to simplify the discussion of the structure. Assuming that each Attribution introduces its own sub-graph, then Figure 6 contains three sub-graphs relevant to our current discussion: (i) the graph of the discourse preceding the first report, which in this case includes only χ , (ii) the sub-graph for the first report, namely Attribution(α, β) and (iii) the sub-graph for the second report, namely Attribution(γ, δ).

If we allow that either part of an Attribution can attach to the Right Frontier of the incoming graph, then the attachments in Figure 6 (and Figures 2, 3, and 4) become unproblematic. The Right Frontier for the sub-graph Attribution(α, β), i.e. the set of nodes to which α and β can be attached, is the singleton set $\{\chi\}$. β attaches to χ , so this attachment is unproblematic. Next, the Right Frontier for Attribution(γ, δ) is $\{\chi, \alpha, \beta\}$ because β is Last and both α and χ directly dominate β through subordinating relations (indicated by downward, as opposed to left-right, arrows). Both γ and δ attach to nodes in this Right Frontier, namely α and χ , respectively.²⁸

By loosening the Right Frontier Constraint in the way suggested for Attributions, we achieve something that two-dimensional accounts aim to achieve: a separation of the attribution content of a discourse parenthetical report from its embedded content. This distinction seems motivated after all—in order to use a discourse parenthetical report, a speaker must first set up someone else’s commitments before she can use them. At the same time, my proposal captures this distinction without appealing to multiple dimensions of meaning and it allows the attribution predicates of a

28. This proposed amendment to the Right Frontier Constraint is slightly complicated by (16), in which Last for the report in (16d) is (16c). The Right Frontier for (16d) is the singleton set $\{(16c)\}$, yet the embedded clause of (16d) attaches to (16a) via \diamond Explanation, suggesting that my adjustment of the Right Frontier fails in this case. I propose, however, that this can only happen because *semantically*, (16c) and (16d) describe the same eventuality so the Right Frontier for (16d) is in effect the Right Frontier for (16c), which includes (16a) and both parts of the report in (16b). Had (16d) been attached to the preceding unit via any discourse relation other than Elaboration, there could not have been a connection back to (16a). Consider a variant of (16) with an Explanation instead of Elaboration: “The leaders of the earthquake investigation disagree because scientists say/have said that the quake was the result of continued oil and gas exploration...”. In this variant, I would argue that potential explanation between (16a) and the embedded clause of the report is not salient; the focus is rather on the leaders of the earthquake and we should expect the discourse to continue with more discussion of them before, if ever, turning to a discussion of what the scientists said. It would be interesting to make an empirical study of continuations of reports and how these bear on the discourse parenthetical/non-parenthetical distinction.

discourse parenthetical report to play the same kinds of discourse functions as the embedded clauses do, which does justice to the observation of *distinct functions*.

5. Syntactic parentheticals

The main goal of this paper has been to develop a semantic account of the discourse function of discourse parenthetical reports, which have a non-parenthetical syntax. Before concluding this discussion, however, this section takes a brief look at how syntactic parentheticals such as (25) might figure in the picture developed so far.

(25) He was out of town, Jill said.

Following Ross (1973) and Simons (2007), I will call the main clause, i.e. *he was out of town* in (25), the *slifted clause*, and I will call the syntactically parenthetical clause, i.e. *Jill said* in (25), the *parenthetical clause*.

Syntactic parentheticals resemble their discourse counterparts in that they can be used when a speaker is not fully committed to the content of the slifted clause (more on that below). Moreover, it is even clearer with syntactic parentheticals that the slifted clause conveys the main point of the utterance and the parenthetical serves an evidential or source-providing function. In addition, we expect that the slifted clause of a syntactic parenthetical will attach to the preceding discourse (when not discourse-initial) just like a discourse parenthetical. Unlike discourse parentheticals, however, parenthetical clauses in syntactic parentheticals out-scope logical operators²⁹

(26) Maybe John was out of town, Jill said.

(27) Maybe Jill said that John was out of town.

(26) does not receive the same reading as (27), in which the attribution predicate is itself in the scope of the possibility modal. Furthermore, syntactic parentheticals do not generally accept negated parenthetical clauses (de Cornulier, 1978; Maier and Bary, 2015):

(28) ?? John was out of town, Jill didn't say.

(29) Jill didn't say that he was out of town.

Their discourse contributions also seem to be different. Compare (30) to (19).

(30) He was out of town, Jane said, but he's just avoiding me, Bill said.

This continuation is clearly more marked than that in (19).³⁰ Simons (2007) provides further discussion of differences between discourse and syntactic parenthetical clauses. I follow Simons in concluding that discourse and syntactic parentheticals really do have a different underlying syntactic structure.

29. I do not know who first noticed that the parenthetical clause of a syntactic parenthetical can outscope logical operators, but it is widely assumed now in discussions of at-issue and not-at-issue content that parenthetical clauses are not-at-issue and that modal operators do not target not-at-issue content. See Potts (2005) and Simons et al. (2010) for recent discussions.

30. There seems to be a difference between reports that appeal to a speech event and reports that merely provide a source for a point of view. Compare: He was out of town, according to Jane (/says Jane), but he's just avoiding me, according to Bill (/says Bill).

So what shall we make of their contribution to *discourse* structure? In the newspaper articles that I have reviewed in studying parenthetical reports, I can isolate three general cases in which syntactic parentheticals tend to be used. First, they are frequently used to extend report contexts. In (31), taken from an article about the Zika virus, the author introduces a spokeswoman for an association of cruise ships who is addressing concerns for travellers heading to areas infected with Zika, which is carried by mosquitos.

- (31) **An association spokeswoman, Elinore Boeke, said** travelers should check with public health officials about destinations they planned to visit and cited the C.D.C.’s current travel advisories, which suggest only that all visitors avoid mosquito bites by using repellent and long clothes.

Cruise ships publish daily fliers on health and safety and instruct passengers on how to avoid bites, **Ms. Boeke added**.³¹

In this example, the syntactic parenthetical *Ms. Boeke added* is used to indicate that the author is extending a report that was begun in the preceding paragraph, and the choice to use a syntactic parenthetical in this case seems to be a stylistic one. Examples like (31) are interestingly similar to Bary and Maier (2014)’s examples of *Unembedded Indirect Discourse* involving the optative in Ancient Greek and the subjunctive in German (following the work of Fabricius-Hansen and Sæbø (2004)).

Syntactic parentheticals are also frequently used with direct quotation. (32) is taken from an article about a report on protective measures taken to protect nuclear facilities in twenty countries. The first paragraph claims that efforts to increase protection of U.S. nuclear power have slowed considerably and the quote at the beginning of the second paragraph elaborates on this judgment.

- (32) The report also concludes that President Obama’s global initiative to sweep up loose nuclear material has slowed substantially.

“There was great progress for six or so years,” **said Sam Nunn**, a former senator who spearheaded efforts to dismantle nuclear weaponry ... and [became] a founder of the Nuclear Threat Initiative. “But it has slowed down. It’s hard to keep this subject on the front burner”.³²

Finally, syntactic parentheticals are commonly used, in newspaper texts at least, to express factual reports in which full author commitment to the content of the slifted clause is inferred. (33), an illustration of a factual report, is taken from an article on the Ebola virus, the most recent outbreak of which was coming to an end at the time the article was written. The excerpt below warns that although there are few to no cases left, vigilance is still needed to prevent further outbreaks.

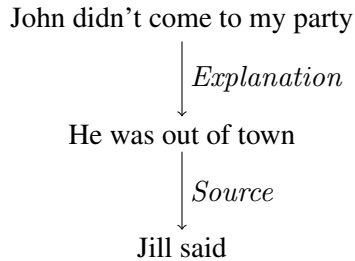
- (33) But in a statement released in Geneva, Ms. Chan added that “our work is not done, and vigilance is needed to prevent new outbreaks.”

31. “C.D.C. May Warn Pregnant Women Against Travel to Countries With Zika Virus,” by D. G. McNeil Jr. *The New York Times*, January 13, 2016.

32. “Nuclear Facilities in 20 Countries May Be Easy Targets for Cyberattacks,” By D. E. Sangerjan. *The New York Times*, January, 14, 2016

The immediate threat stems from persistence of the virus in body fluids, notably in the semen of male survivors, up to a year after they are free of the disease and show no symptoms, **said Rick Brennan**.³³

To handle these data, I propose that syntactic parentheticals contribute a structure like that posited for Source in Hunter et al. (2006). That is, the parenthetical clause is subordinate to the slifted clause, which connects directly to the preceding discourse, without triggering a modal discourse relation. Thus, we get an annotation like that below for (25) as a continuation of (1a):



Source is a veridical relation, meaning that it entails both of its arguments. This therefore straightforwardly handles the factual uses of syntactic parentheticals as in (33).

Modelling extension cases, e.g. (31), is slightly more complicated, but they should not pose a problem for Source, because the slifted clause should be in the (discourse) scope of an Attribution. The role of the parenthetical clause seems to be to extend the scope of an Attribution so that the slifted clause is combined with material in the scope of a previous Attribution to form a *complex discourse unit*. The parenthetical clause would not figure in this complex discourse unit, of course, but might function more as a presupposition that ties the parenthetical report to a report introduced previously in the context (cf. Bary and Maier (2014) and Fabricius-Hansen and Sæbø (2004)). In this case, the veridicality requirement enforced by Source would only need to be satisfied inside of the main Attribution's scope; the author of the main discourse need not commit to the content of the slifted clause.

An analysis of direct quotation cases such as (32) would need to appeal to a theory of quotation, and it is not in my plans to outline a solution here or to choose an appropriate theory of quotation. I merely note that the fact that the quotation marks are needed to distance the main author from the report agent means that the quotation marks are themselves signalling some kind of special context. The parenthetical clause on its own cannot achieve this, so the quotation marks (or whatever signal might be given orally) are signalling something in addition.

The hypothesis sketched here entails that a speaker must commit to the slifted clause of a syntactically parenthetical report unless the report extends another speech context.³⁴ Thus, I would

33. "Deadliest Ebola Outbreak on Record Is Over, W.H.O. Says," by D. Searcey, N. Cumming-Bruce and C. MacDougall. *The New York Times*, January 14, 2016.

34. This claim requires two caveats. First, I suspect that fiction is more liberal with regard to constraints on indirect reports, but my focus in this paper is on how reports are used in non-fictional contexts and in newspaper articles in particular. (See Maier and Bary (2015), and their example (19), for discussion of the difference between narrative and non-fiction environments.) Second, there seems to be a difference between first-person reports and third-person reports that is relevant to syntactic parentheticals. Compare the continuations (i) and (ii):

† John didn't come to my party

(i) He was out of town, I suspect.

not expect a speaker to start a conversation with (34) or to use it to answer a question like *Where is John?* or *Why didn't John come to your party?*.

(34) John/He is out of town, Mary said, but I bet he's not.³⁵

This is intuitively correct. A way to reconstruct what is going in (34) is to say that Mary says p but the speaker says $\neg p$; fundamentally, this is a “he said, she said” style of Contrast like the Contrast that we see in (19). As the example stands, however, it has the structure of “ q but I bet that p ”, where q is *John is out of town* and *I bet that p* turns out to be equivalent to *but I bet that $\neg q$* . This is not a good Contrast because it implies that the speaker is disagreeing with herself. To achieve the appropriate Contrast, the two arguments for Contrast need to be structurally parallel (Asher and Lascarides (2003)), as in *Mary said that John is out of town but I bet he's not*. Furthermore, because commitments are being called into question in (34), the content of the parenthetical clause, which is crucial for securing these commitments, should play a leading role, not the backgrounding role that it plays in (34).

Modelling syntactic parentheticals with Source would naturally explain the infelicity of examples like (28). Normally, a speaker would use a report like (29) when she does *not* have evidence that John is out of town (and expects that Jill would have provided such evidence had such evidence been available). If Source is a veridical relation, then we predict that this reading would be blocked for (28) because Source would require the speaker to commit to the truth of John's being out of town—the exact opposite of what happens in (29). The only other use of (29) that I can imagine is in Corrections, e.g. *Jill didn't say he was out of town, she said he was in a meeting*. But in such Corrections, the crucial issue is precisely what Jill said, so we shouldn't expect a parenthetical structure, which would place this issue in a less prominent position, to be chosen.

The hypothesis developed in this section for syntactic parentheticals would be disproven by clear evidence that speakers regularly use syntactically parenthetical constructions to convey indirect reports outside of the contexts illustrated by (31) and (33). This evidence is lacking in the data sets that I have used, but lack of evidence in my corpus is far from conclusive support for my analysis. Note, however, that a collection of syntactically parenthetical reports taken in isolation of context would *not* undermine my hypothesis, even if each report was judged felicitous by an appropriate set of native speakers. Respondents might judge such a report to be acceptable because he or she could imagine a context in which it would be acceptable. What is needed to disprove my hypothesis is genuine evidence that speakers actually use or would use these reports in contexts other than the three that I have described.

(ii) He was out of town, Jill suspects.

I have not surveyed native speakers on these continuations, but my personal impression is that (ii) is marked. Certainly some empirical work would need to be done to test this judgment; however, because the empirical work that I have been able to carry out pertains almost exclusively to third-person reports, the difference between first and third-person reports is a topic that I will not explore in this paper.

35. An example like: *John is out of town, Mary said, although she could be wrong* would be better, at least as long as there is a pause just before *although* that signals that the clause is a kind of afterthought. I would analyze such examples as Corrections: the speaker first makes a strong claim and then weakens it with the subordinate clause. If this is a Correction, however, it does not require a revision of the hypothesis that I have presented here.

6. A new look at discourse parenthetical reports

The aim of this paper has been to make concrete the generally accepted idea that a discourse parenthetical reading of a (syntactically non-parenthetical) English speech report results from the way that the report is *used* in a discourse, rather than from the syntax or semantics of the report alone. The account that I have offered extends tools developed by rhetorical theories, whose main purpose is precisely to capture how relations between discourse units give rise to semantically relevant content in the discourse that cannot be inferred from the semantics of the units alone. Rhetorical theories suggest a particular way of understanding what it means for the embedded clause of a discourse parenthetical report to play a special discourse function: it means that the embedded clause can, independently of the attribution predicate, enter into discourse relations with units that precede the report and lie outside of another report context. This is what it is for a report to be discourse parenthetical in my account. However, applying this notion to discourse parenthetical reports has required that we extend rhetorical theories in certain ways, most notably by requiring that discourse parenthetical reports trigger modal discourse relations and by modifying the Right Frontier Constraint to capture these connections and the fact that both parts of a discourse parenthetical report can be rhetorically relevant, as illustrated by (16), (19), and (20). The resulting account is an attachment-based solution, but one that overcomes the problems of *hedged commitments*, *no commitment* and *distinct functions*.

Analyzing reports in this way answers some outstanding questions. First, because of the semantics of veridical relations, the use of a discourse parenthetical report entails commitment to the possibility that the embedded content is true. Normally, this commitment is carried by the speaker of the discourse, though in the case of reports embedded in larger speech contexts, the commitment can be carried instead by the agent of that context. Note that this consequence does not entail that a report of the form ‘X said that *p*’ semantically entails $\diamond p$, which would be too strong for non-parenthetical readings of such reports; the commitment to $\diamond p$ merely results from the fact that a discourse parenthetical report contributes a relation $\diamond R$ to the discourse structure. Second, and conversely, a report is not used parenthetically if the speaker explicitly denies the content of the embedded clause; her denial rather serves to cancel the inference of a relation. This resolves the tension in Simons (2007) regarding cases of *no commitment* (see §3.2 of this paper): on the one hand, examples of *no commitment* seem discourse parenthetical because we would have treated them as such had the speaker not gone on to deny the embedded content. On the other hand, a speaker doesn’t actually use the embedded contents of these reports to answer a question or explain an eventuality, etc. My account treats the latter concern as primary; we’re out to track what inferences an interpreter can draw from a discourse, so we need to look at how the reports are actually used.

My account also has consequences for current debates about at-issue or main-point content (Potts (2005), Simons et al. (2010), Murray (2014)). Previous accounts of discourse parenthetical reports, including Simons’ question-based suggestion and attachment-based approaches in rhetorical theories, presuppose implicitly that a report serves to accomplish a single function in a discourse and then try to determine what that function is. What examples like (16), (19), and (20) show is not only do the attribution predicate and embedded clause of a discourse parenthetical report work together to determine a single attachment to the preceding discourse; sometimes, they determine two independent attachments whose relation to one another can only be inferred by looking at the larger discourse context. In Simons’ terminology, sometimes parenthetical reports are used to simultaneously answer two different questions (neither of which need be a sub-question of the other). Similarly, examples of speaker disagreement and strengthening show that the import of a

report—what answer it provides, if you will—may not be clarified by looking only at its immediate attachment in the discourse. Standing back and looking at its role in the larger discourse might give us an entirely different answer. But then, how do we assess the at-issue status of that particular discourse move? The complications that arise in studying reports in discourse show that we need to be careful about how to frame the debate. Further investigation of the relation between at-issueness and discourse structure will have to be saved for another paper, however (though see Asher (2000), Hunter and Asher (forthcoming), and Jasinskaja (2016) for some relevant discussion).

References

- Nicholas Asher. *Reference to Abstract Objects in Discourse*. Number 50 in Studies in Linguistics and Philosophy. Kluwer, Dordrecht, 1993.
- Nicholas Asher. Truth conditional discourse semantics for parentheticals. *Journal of Semantics*, 17: 31–50, 2000.
- Nicholas Asher and Alex Lascarides. *The Logics of Conversation*. Cambridge University Press, 2003.
- Corien Bary and Emar Maier. Unembedded indirect discourse. In Urtzi Etxeberria, Anamaria Fălăuş, Aritz Irurtzun, and Bryan Leferman, editors, *Proceedings of Sinn und Bedeutung 18*, pages 77–94, 2014.
- Matthias Buch-Kromann and Iørn Korzen. The unified annotation of syntax and discourse in the copenhagen dependency treebanks. In *the ACL Linguistic Annotation Workshop*, 2010. URL <http://dl.acm.org/citation.cfm?id=1868737>.
- Matthias Buch-Kromann, Daniel Hardt, and Iørn Korzen. Syntax-centered and semantics-centered views of discourse. can they be reconciled? In Stefanie Dipper and Heike Zinsmeister, editors, *Beyond Semantics: Corpus-based Investigations of Pragmatic and Discourse Phenomena*, volume 3, pages 17–30, Göttingen, 2011.
- Lynn Carlson and Daniel Marcu. Discourse tagging reference manual. Technical report, Information Sciences Institute, 2001. URL <ftp://128.9.176.20/isi-pubs/tr-545.pdf>. ISI Technical Report ISI-TR-545.
- Laurence Danlos and Owen Rambow. Discourse relations and propositional attitudes. In *Proceedings of the Constraints in Discourse Workshop (CID 2011)*, Agay, France, 2011. URL <http://hal.inria.fr/inria-00614763>.
- Benoît de Cornulier. L’incise, la classe des verbes parenthétiques et le signe mimique. *Cahier de linguistique*, 8:53–95, 1978.
- Marie-Catherine de Marneffe, Christopher D. Manning, and Christopher Potts. Did it happen? the pragmatic complexity of veridicality assessment. *Computational Linguistics*, 38(2):301–333, 2012.
- Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Rashmi Prasad, and Aravind Joshi. 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, 2005. URL <http://www.aclweb.org/anthology/W/W05/W05-03.pdf>.
- Cathrine Fabricius-Hansen and Kjell Johan Sæbø. In a mediative mood: The semantics of the german reportive subjunctive. *Natural Language Semantics*, 12:213–257, 2004.
- Jonathan Ginzburg. *The interactive stance*. Oxford University Press, 2012.
- Jeroen Groenendijk and Martin Stokhof. *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, University of Amsterdam, 1984.
- Daniel Hardt. A uniform syntax and discourse structure: the copenhagen dependency treebanks. *Discourse and Dialogue*, 4(2), 2013. URL <http://elanguage.net/journals/dad/article/view/2830>.
- Joan B. Hooper. On assertive predicates. In J.P. Kimball, editor, *Syntax and Semantics*, volume 4, pages 91–124. Academic Press, NY, 1975.
- Julie Hunter and Nicholas Asher. Shapes of conversation and at-issue content. In *Proceedings of Semantics and Linguistic Theory 26*, forthcoming.
- Julie Hunter, Nicholas Asher, Brian Reese, and Pascal Denis. Evidentiality and intensionality: Two uses of reportative constructions in discourse. In *Proceedings of the Constraints in Discourse Workshop (CID 2006)*, Maynoth, Ireland, 2006.
- Katja Jasinskaja. Not at issue any more. Unpublished manuscript presented at the 17th Szklarska Poręba Workshop on the Roots of Pragmasemantics, 2016.
- Hans Kamp and Uwe Reyle. *From Discourse to Logic*. Kluwer Academic Publishers, 1993.
- Emar Maier and Corien Bary. Three puzzles about negation in non-canonical speech reports. In *Proceedings of the 20th Amsterdam Colloquium*, 2015.
- William Mann and Sandy Thompson. Rhetorical structure theory: Towards a functional theory of text organization. *Text*, 8:243–281, 1988. URL http://www.sfu.ca/rst/05bibliographies/bibs/Mann_Thompson_1988.pdf.
- Christian Matthiessen and Sandra A. Thompson. *The Structure of Discourse and ‘Subordination’*, pages 275–329. John Benjamins Publishing Company, 1988.
- Sarah Murray. Varieties of update. *Semantics and Pragmatics*, 7(2):1–53, March 2014. doi: 10.3765/sp.7.2.
- Christopher Potts. *The Logic of Conventional Implicatures*. Oxford Studies in Theoretical Linguistics. Oxford University Press, Oxford, 2005.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Aravind Joshi, and Bonnie Webber. Attribution and its annotation in the Penn Discourse Treebank. *Revue TAL*, 47(2), 2007a.

- Rashmi Prasad, Eleni Miltsakaki, Nikhil Dinesh, Alan Lee, Aravind Joshi, Livio Robaldo, and Bonnie Webber. The Penn Discourse Treebank 2.0 annotation manual. Technical report, Institute for Research in Cognitive Science, University of Philadelphia, 2007b. URL http://repository.upenn.edu/ircs/_reports/203/.
- Gisela Redeker and Markus Egg. Says who? on the treatment of speech attributions in discourse structure. In *Proceedings of the Workshop Constraints in Discourse*, pages 140–146, 2006.
- Brian Reese, Julie Hunter, Nicholas Asher, Pascal Denis, and Jason Baldrige. Reference manual for the analysis and annotation of rhetorical structure. Technical report, The University of Texas at Austin, Austin, Texas, 2007.
- Craige Roberts. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, 5:1–69, 2012. doi: <http://dx.doi.org/10.3765/sp.5.6>.
- Johan Rooryck. Evidentiality, parts i and ii. In *Glott International* 5, 4 and 5. Blackwell, 2001.
- John Ross. Slifting. In Maurice Gross, Morris Halle, and Marcel Schutzenberger, editors, *The formal analysis of natural languages*, pages 133–169, The Hague, 1973. Mouton.
- Roser Saurí and James Pustejovsky. Factbank: A corpus annotated with event factuality. In *Language Resources and Evaluation*, volume 43, pages 227–268, 2009.
- Mandy Simons. Observations on embedding verbs, evidentiality, and presupposition. *Lingua*, 117 (6):1034–1056, 2007.
- Mandy Simons, Judith Tonhauser, David Beaver, and Craige Roberts. What projects and why. In David Lutz, editor, *Proceedings of SALT XX*, 2010.
- James Opie Urmson. Parenthetical verbs. *Mind*, 61 (244):480–496, 1952.
- Antoine Venant. *Pristine Perspectives on Logic, Language, and Computation: ESSLLI 2012 and ESSLLI 2013 Student Sessions. Selected Papers*, chapter Toward a Discourse Structure Account of Speech and Attitude Reports, pages 233–246. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- Antoine Venant and Nicholas Asher. Ok or not ok? commitments in acknowledgments and corrections. In *Semantics and Linguistic Theory*, volume 25, pages 595–614, Stanford, 2016. doi: 10.3765/salt.v25i0.3072. URL <http://journals.linguisticsociety.org/proceedings/index.php/SALT/article/view/25.595>.
- Antoine Venant, Nicholas Asher, and Cédric Dégrement. Credibility and its attacks. In *Workshop on the Semantics and Pragmatics of Dialogue - SemDIAL 2014*, pages pp. 154–162, Edinburgh, United Kingdom, 2014. URL <https://hal.archives-ouvertes.fr/hal-01145832>.
- Bonnie Webber. D-Itag: extending lexicalized tag to discourse. *Cognitive Science*, 28:751–779, 2004.