

Attachment in Syntax and Discourse: Towards an explanation for the flexible scope of non-restrictive relative clauses*

Katja Jasinskaja
University of Cologne

Claudia Poschmann
University of Cologne
University of Frankfurt a.M.

Abstract Schlenker (2013) gives a number of puzzling counterexamples to the widely accepted claim that non-restrictive relative clauses (NRCs) are always interpreted with respect to the global context, and never in the scope of entailment-canceling operators such as *if*. Local readings are available for NRCs attached to their host clause by a coordinating coherence relation. This paper develops a theoretical explanation of this pattern. We argue that NRCs are interpreted locally only if they are attached locally to their host clause both in syntax and in discourse structure. Subordinating coherence relations like *Elaboration* and *Explanation* resist discourse attachment in the scope of *if* because they tend to go together with relations that can only hold between speech acts. Like other subordinate clauses, NRCs tend to express subordinating coherence relations, which ultimately explains their pervasive tendency for global interpretation. In other words, this study shows how a theory of discourse coherence can help solve a problem in sentence semantics.

Keywords: projection, appositives, relative clauses, coherence relations

1 Introduction

As is well known, non-restrictive relative clauses (NRCs) present a problem for compositional semantics. Despite their embedded position, they typically do not contribute to the meaning of the embedding sentence. Rather, they are quite stubbornly interpreted like independent sentences. For example, the sentence in (1a) entails that the dean hates the speaker, although the NRC that provides this information is embedded in the antecedent of a conditional.

* The research reported in this paper was funded by the German Research Foundation (DFG) as part of the SFB 1252 *Prominence in Language* in the project C06 *Prominence in subordinating rhetorical relations* at the University of Cologne. We also thank the audiences of SALT and of the Workshop on *At-issueness, scope and coherence* (ASC, July 9-10, University of Cologne) as well as the anonymous reviewers of SALT for their insightful comments on previous versions of this paper.

- (1) a. If Peter calls the dean, who hates me, I would be in deep trouble.
 b. \rightsquigarrow The dean hates me.
 c. $\not\rightarrow$ If Peter calls the dean and the dean hates me, I would be in trouble.

Standard accounts are designed to ensure invariantly a wide-scope interpretation of NRCs. Basically, there are two ways to do this: Syntactic approaches assume that NRCs outscope the embedding clause because they are attached high in the syntactic tree (at CP-level, McCawley 1982), or even syntactically orphaned (Safir 1986; Fabb 1990; Espinal 1991). Alternatively, it has been argued that NRCs are attached locally to their host clause at DP (Del Gobbo 2003; Potts 2005) or at IP level, but that the content they contribute *projects*, i.e. is interpreted globally, outside the context of the embedding sentence. This happens because that content is semantically (Potts 2005) or pragmatically (AnderBois, Brasoveanu & Henderson 2011; Simons, Tonhauser, Beaver & Roberts 2011) not-at-issue, or because NRCs have their own illocutionary force (Koev 2013).

Recent work on NRCs, however, has challenged this view showing that the content of NRCs interacts in multiple ways with the content of the main clause (Nouwen 2007; Arnold 2007; AnderBois et al. 2011; Sæbø 2011). Moreover, Schlenker (2013: 7) points out English and French examples where the NRC is interpreted locally in the scope of entailment-canceling operators like *if*. For instance, (2a), adapted from Schlenker 2013: 7, does not tell us that the dean called the chair, but instead has the local reading in (2c): The NRC contributes to the restriction of the conditional.

- (2) a. If Peter called the dean, who then called the chair, I would be in trouble.
 b. $\not\rightarrow$ The dean called the chair.
 c. \rightsquigarrow If P. called the dean and the dean called the chair, I would be in trouble.

First experimental results from German confirm the existence of such local readings, but suggest that their availability is dependent on the type of coherence relation established between the NRC and its host clause (Poschmann 2018). These findings challenge existing approaches to NRCs. No worked-out explanation of this phenomenon has been offered so far. The present paper addresses the following questions: How can we explain the existence of local readings of NRCs? Why do coherence relations affect the scope of NRCs? And why are global readings still preferred most of the time?

After a brief empirical overview in section 2, we argue that NRCs are locally attached to their host clause in the syntax, but not necessarily semantically conjoined in that position. Rather they are connected to their host clause by an underspecified coherence relation which must be resolved pragmatically (section 3). That pragmatic

inference is guided by the same principles as the inference of coherence relations between independent sentences in discourse. In section 4, we show how the traditional discourse-theoretic distinctions between subordinating and coordinating relations, and between semantic and presentational relations help explain the role played by coherence relations in licensing local readings of NRCs. Finally, section 5 addresses the question of why NRCs project most of the time.

2 Local Readings of NRCs

One of the features that can reveal the position in which an NRC is interpreted is its tense. The past tense of the NRC in (2) is an instance of ‘fake tense’, that is, it is bound by the conditional and does not indicate an event in the past. As Schlenker (2013: 7) points out, replacing the NRC by a parenthetical (3) or a postposed matrix clause (4) makes the sentence less acceptable because the past tense is no longer bound by the *if*-clause. This shows that the NRC in (2) is indeed interpreted locally.

- (3) ?? If Peter called the dean (she then called the chair), I would be in trouble.
(4) If Peter called the dean, I would be in trouble. ?? She then called the chair.

However, NRCs do not always allow for local readings. If we try to force the NRC in (1) into the scope of the conditional by changing its tense to past, the sentence becomes awkward (5). The past tense seems to resist the bound reading.

- (5) ?? If Peter called the Dean, who hated me, I would be in trouble.

Apparently, NRCs allow for local readings only under some special conditions. But which are those conditions? What makes the difference between (1) and (2)?

Apart from tense, the examples differ, for instance, in the presence/absence of the anaphoric adverbial *then*. A straightforward explanation, suggested by Martin (2016), relates the local interpretation in (2) to the need to bind the anaphoric adverbial *then*, which only finds an appropriate antecedent (the event of Max calling the dean) in the scope of the conditional—a solution parallel to van der Sandt’s (1992) trapping of presuppositions. However, one should then expect that the pronoun *him* in (6), which can only be bound by *someone* under *if*, should also be able to trap the NRC. But, unfortunately, the sentence receives no coherent interpretation, although the corresponding conjoined conditional does.

- (6) If someone_a wears this jacket, and it / #which is too big for him_a, he will look silly.

Extending the notion of anaphora that operates in presupposition trapping from identity to (weakly) causal relations, proposed by Zeevat (2016), does not really help

either. Even if the NRC in (2) might seem to be trapped by a causal relation—the dean would not call the chair if Peter did not call the dean—(7) shows that an NRC can receive an embedded interpretation even in the absence of a causal link or with inverted temporal order.

- (7) If Mary stands in front of Peter, who (first/then/also) moves a little closer to Max, everyone will fit into the picture.

Schlenker (2013: 42–44) suggests a more promising generalization: The availability of local readings depends on the coherence relation between the NRC and its host clause. In (8) we list some common definitions of coherence relations that we assume throughout this paper (Jasinskaja & Karagjosova forthcoming), which are not substantially different from those used by Schlenker (2013).

- (8) *Elaboration*(α, β): α and β describe the same eventuality (or entity, in case of *Entity-Elaboration*)
Explanation(α, β): β describes the cause of the eventuality in α
Result(α, β): α describes the cause of the eventuality in β
Narration(α, β): β describes an eventuality that follows that of α in time and space, but is not necessarily caused by it
Parallel(α, β): α and β describe similar eventualities
Contrast(α, β): α and β describe ‘opposite’ eventualities, that is, α non-monotonically implicates p , while β implicates $\neg p$

Schlenker suggests that *Narration* and *Result* holding between the NRC and its host clause, cf. (2), make local readings of the NRC possible, whereas *Elaboration* (9) and especially *Explanation* (10) resist local interpretation most strongly.

- (9) If John engaged his lawyer, who his mother ordered to file a complaint, ...
 \rightsquigarrow John’s mother ordered his lawyer to file a complaint.
 $? \rightsquigarrow$ If John’s mother ordered his lawyer to file a complaint and thereby John engaged him/her, ...
- (10) If John called his aunt, who was / ??were a lawyer, ...
 \rightsquigarrow John’s aunt used to be a lawyer.
 $\not\rightsquigarrow$ If John called his aunt because she is a lawyer, ...

Interestingly, this division corresponds to another fundamental division in discourse structure—that between *coordinating* (*Narration, Result, Parallel, Contrast*) and *subordinating* coherence relations (*Elaboration, Explanation*) in the sense of Asher & Vieu 2005, which we will discuss in more detail later. The hypothesis that discourse-structural coordination/subordination distinction lies at the core of the phenomenon at hand is investigated by Poschmann (2018). In a set of acceptability

tests, she compares German NRCs under *wenn* ‘if’ to conjoined and parenthetical clauses, like the English (3), as well as to postposed matrix clauses, like (4). Further variables in the study are the predicate type (state/event) and the presence/absence of discourse adverbials such as *dann* ‘then’. In a context that supports local interpretation, NRCs turn out to be less acceptable than conjoined clauses (about 50% vs. 91%, respectively), but significantly more acceptable than parentheticals (21%) and postposed matrix clauses (9%). Moreover, forcing a coordinating coherence relation by discourse adverbials that signal *Narration* or *Contrast* raises the acceptability of local readings of NRCs to about 83% regardless of predicate type. These results corroborate Schlenker’s intuitions and support Poschmann’s stronger version of the generalization. Further support comes from examples like (7), where the local interpretation of the NRC appears to be licensed by the *Parallel* relation to the conditional clause. In other words, a whole variety of coordinating coherence relations can hold between the NRC and its host clause embedded under entailment-canceling operators like *if*, whereas subordinating relations tend to require global interpretation.

This looks like a nice empirical generalization. However, even if we ignore the somewhat less clear status of *Elaboration* (9) for the time being, the pattern still needs to be explained. Why should coherence relations affect the scope of NRCs in the first place? Why is the distinction between discourse-structural coordination and subordination relevant? And finally, why are NRCs interpreted globally most of the time despite the variety of possibilities for local interpretation that we have discovered? These are the questions we address in the rest of this paper.


3 Attachment in syntax and semantic interpretation

Before we explain how different kinds of coherence relations affect the projection pattern of NRCs, we should clarify our assumptions about the syntactic and semantic representation that serves as input to the inference of coherence relations. What is the position of NRCs in the syntactic structure of a sentence? How are they semantically interpreted? And how are they similar to or different from clauses of other kinds?


We follow Schlenker (2013) and Poschmann (2018) in assuming that an NRC contributes a proposition of type $\langle t \rangle$ (after having been coindexed with the referent of its host DP) and is adjoined not higher than to the IP node of the host clause and not lower than to the node of the host DP. Higher attachment would be problematic for an account of the bound readings of tense in NRCs discussed in section 2. A lower position, by contrast, would lead to an interpretation in which the relative clause restricts the denotation of the DP (Heim & Kratzer 1998). We remain impartial regarding the question of the exact locus of NRC attachment: whether it is the IP (as, for instance, in Schlenker’s proposal), the DP (currently the standard assumption taken in most accounts of NRCs, e.g. Del Gobbo 2003; Koev 2013; Poschmann

2018), or anywhere in between. We also disregard the question whether the NRC denotes a predicate $\langle et \rangle$ like its restrictive counterpart (Heim & Kratzer 1998; Koev 2013; Schlenker 2013) or a proposition $\langle t \rangle$. Since in $\langle et \rangle$ accounts the NRC is coindexed with its host DP anyway, the contribution of the NRC in the end is still a proposition. Thus, whichever approach we assume, we get two propositions at the IP level: that of the host IP and that of the NRC.

One of the central claims of our proposal is that there is no semantic rule that simply allows to conjoin adjacent propositions, unless there is an actual explicit connective present, such as *and* or *but*, which contributes a conjunction as part of its encoded semantic content. Instead, adjacent propositions—those of the host clause and the NRC in particular—are treated much like juxtaposed sentences in discourse. They are connected by an underspecified coherence relation **R**, as shown in (11), which has to be resolved by pragmatic reasoning.¹




This inference process may or may not lead to success. If it does, **R** is instantiated by a specific coherence relation, e.g. *Narration* in (2). The proposition that results from connecting the NRC and the host IP by *Narration* is assigned to the higher IP, and if the latter is in the scope of an entailment-canceling operator, the entire *Narration* including the content of the NRC is interpreted in the scope of that operator, cf. (12). Since *Narration* subsumes the semantics of the natural language conjunction *and* (more on this in section 4.3), the inference leads ultimately to conjoining the propositions of the NRC and the host IP in the scope of *if*. This is how we get local readings as in (2). Thinking of coherence relations as a kind of pragmatic inference, this is similar to the phenomenon of local quantity implicature (Chierchia 2004) intruding into the compositional semantics of a sentence.



¹ A similar idea is put forward in Schlenker 2013: 20, who argues that NRCs are semantically connected to their host clause by an operator • described as ‘the operation by which parataxis in discourse is evaluated’. For a fully spelled out analysis see Poschmann 2018.

However, more often than not we fail to establish a coherence relation between the content of an NRC and that of its host clause at the local level. In that case, only the proposition of the host clause serves as input to further semantic composition, the relation \mathbf{R} to the proposition expressed by the NRC remains underspecified and keeps being percolated to higher levels, until it finds a higher discourse attachment site, which will normally be located at the root clause level, cf. (13).² This leads to more generally familiar global, or projective readings of NRCs, as in (1).



In sum, an NRC is interpreted locally in the scope of an entailment-canceling operator O only if it is attached to a clause in the scope of O both syntactically and discourse-structurally. Assuming that syntactic attachment of NRCs is local, all depends on the result of integrating the NRC in the discourse structure. The difference between NRCs (2) and clausal parentheticals (3) consists in the assumption that the latter are not syntactically related to their host clause. Therefore, there is no pressure to integrate parentheticals in the semantic composition, they are interpreted as root clauses right away and therefore only have global readings. The property of being syntactically integrated, in turn, is shared by NRCs and coordinate clauses:

(14) If Peter called the dean *and* she called the chair, I would be in trouble.

However, in coordinate clauses the connective *and*, *but* or *or* tells us explicitly how to make one proposition out of two. Those propositions may or may not be enriched with additional pragmatic inferences, i.e. coherence relations in particular, but their semantic integration does not rely on those inferences, a local relation is present by virtue of being linguistically encoded, which guarantees local interpretation. NRC structure by itself tells us *that* the propositions should be semantically related, but not *how* they should be related, and only if the pragmatic inference process leads to the ‘insertion of an implicit *and*’ between the host and the NRC proposition, will NRCs receive local readings similar to those of corresponding conjoined clauses.

4 Coherence relations and the scope of NRCs

It follows from the picture presented above that whether or not an NRC is interpreted locally depends on whether or not it is attached to its host clause by a coherence

² See Schlenker 2013: 17 for discussion of potential intermediate readings.

relation that conjoins the propositions expressed by the clauses. But which factors determine whether this will be the case? In this section we first consider some general characteristics of coherence inferences that determine the attachment point of new discourse material and the choice of coherence relation (section 4.1). We then discuss features of coherence relations that make them particularly suitable or unsuitable for local connection between an NRC and its host clause. The features that we concentrate on are the *source of coherence*, which determines whether a coherence relation can hold in the scope of entailment-canceling operators (section 4.2) and the compatibility of a relation with natural language conjunction (section 4.3). According to these criteria, *Elaboration* or *Explanation* turn out to be largely incapable of local attachment of NRCs, which confirms Schlenker's (2013) observations.

4.1 Inference to the best explanation

We adopt the view going back to Hobbs, Stickel, Appelt & Martin 1993 of discourse interpretation as abduction—a non-monotonic inference mechanism that seeks for the best explanation of the observed utterances. That is, what must the speaker have had in mind for her to produce the given discourse? This puts the inference of coherence relations on a par with Gricean inference (cf. Geurts 2010) as well as more recent approaches to pragmatics such as the Bayesian Rational Speech Act model (Frank & Goodman 2012). The preferred attachment point and type of coherence relation must be likely (relevant, plausible, etc.).

In order to see how much sense local attachment really makes, it is useful to look at the relationship between the propositions expressed by the NRC and its host clause outside the embedding context. As we argued in the previous section, in terms of the degree of implicitness of the coherence relation, NRCs are comparable to juxtaposed independent utterances. To abstract away from the NRC as a specific means of expression and be able to evaluate the relevance and plausibility of the relation as such, one can rephrase the NRC as an independent full sentence or a fragment DP. If the resulting discourse makes sense, then local attachment with the inferred coherence relation is by itself a likely interpretation. Applying this test to (2) and (1) results in (15) and (16), respectively. The discourse in (15) is perfectly coherent and is readily understood as a *Narration*, cf. (8). The adverbial *then* clearly supports this interpretation. In contrast, (16) does not make that much sense without further context. The relation could perhaps be characterized as that between a foreground event and a background state, but it still remains unclear just why the speaker mentions the dean's attitude towards him in connection with the phone call.

(15) Peter called the dean. She then called the chair. *Narration*

(16) ? Peter called the dean. She hates me.

However, relative acceptability is reversed once we include the context of the entire original complex sentence. Now (17) is incomprehensible without further context. The dean's call cannot be the result of Peter's call, because the latter did not take place. Otherwise, what does the call have to do with the speaker's potential trouble? In contrast, (18) makes a lot more sense than (16): The dean's negative attitude towards the speaker definitely increases the probability of trouble after Peter's call, especially if the latter communicates some unfavorable information about him. This suggests that the NRC in (1), is not attached by a discourse relation to its host clause, but rather to the entire sentence: *Explanation*(ϕ , ψ), cf. (19).

(17) ?? If Peter called the Dean, I would be in deep trouble.
She (then) called the Chair.

(18) If Peter called the Dean, I would be in deep trouble. *Explanation*
She hates me.

(19) [If Peter called the Dean, [who hates me] ψ , I would be in deep trouble.] ϕ

Applying abductive reasoning to figure out the NRC's discourse attachment point requires considering possible alternative attachments such as those corresponding to (16) and (18). If the content of the NRC is more relevant in the context of the entire sentence than in the context of the host clause alone and assuming that the NRC is otherwise not a more likely expression for local than for global attachment, the latter ends up as the preferred reading, which places the NRC outside the scope of *if*.³

The availability of a particular alternative attachment site can be affected by its prominence, or the degree of activation at the current point in discourse development. For instance in (20), discourse attachment of the NRC goes together with the resolution of the anaphoric adverbial *then*. By the time the adverbial is interpreted, the temporal location of the event described by the host clause is not yet constructed, whereas the temporal location of the event in the preceding sentence is available and prominent, and is therefore picked out by the adverbial. Again, the NRC is attached outside the scope of *if*, *Narration*(ϕ , χ), and is therefore interpreted globally.

(20) [Nick stuck out his tongue at Jamie.] ϕ adapted from
If [Jamie, [who then hit him,] χ had simply left instead,] ψ Koev 2013
he wouldn't have realized that he insulted her.

³ Note the parallelism between this inference and Gricean relevance implicature. Just as the message *there is a garage around the corner* is more relevant on the assumption that the garage is open rather than closed, the proposition *the dean hates me* is more relevant assuming global attachment than assuming local attachment.

In sum, our main claim is that the pragmatic inference that leads to the integration of the content of an NRC in sentence and discourse semantics follows the same principles and is affected by the same factors as the inference of coherence relations in general. That is, in principle, any coherence relation that can hold between juxtaposed independent sentences can also hold between an NRC and its host clause or broader context. Similarly, whether an NRC is attached to its host clause, the whole sentence, or any other appropriate antecedent, depends on the same factors that determine the attachment site of an independent sentence in discourse structure. This, however, does not explain Schlenker's intuition supported by Poschmann's experiments that NRCs can be attached locally only by a small subset of coherence relations. In the next sections we show that different coherence relations are not equally suitable for local discourse attachment.

4.2 The source of coherence

Whether or not a coherence relation and the discourse segments it connects can be embedded under entailment-canceling operators depends on the *source of coherence* of a given pair of segments (Sanders, Spooren & Noordman 1992), in particular, whether it is the properties of the situations in the world described by those segments or the properties of the speech acts that make the sequence coherent. Obviously, the latter case presupposes that the discourse segments constitute independent speech acts, which cannot be embedded under entailment-canceling operators. Schlenker (2013) observes that NRCs attached by *Elaboration* or *Explanation* resist local interpretation most strongly. In this section we argue that these coherence relations have a tendency to accompany a relation that operates at speech act level and this is what makes them unsuitable for local attachment of NRCs. This will require a brief excursus into theories of discourse coherence.

Coherence of a sequence of sentences can have its source in the relations between states of affairs described by those sentences. The idea is that discourse is coherent because 'the world that is described is perceived as coherent' (Sanders et al. 1992: p. 7). A typical example of a system based entirely on this principle is Kehler 2002. The list of coherence relations in (8) in section 2 *de facto* only contains relations defined in semantic terms, and the same applies to the subsets used by Schlenker (2013) and Poschmann (2018).

Coherence relations characterized in these terms do not impose any relevant constraints on the type of abstract entities they connect (events, propositions vs. speech acts). Causal and spatiotemporal relations can hold between descriptions of eventualities, which can be actual, potential or counterfactual. Therefore, in principle, nothing prevents these relations from holding in the scope of entailment-canceling operators. However, since speech acts are events, they can also be connected by

causal relations, as in the famous example (21) discussed by Sweetser (1990), where the second sentence presents the cause for the speaker's question.

(21) What are you doing tonight? Because there is a nice movie on.

In general, it appears that coherence relations characterized in terms of relations in the world can *happen* to hold between speech acts, because speech acts are also entities in that same world, but the definitions of the relations do not make reference to anything specific to speech acts.

However, coherence may also result specifically from the relationship between speech acts. In particular, one speech act (the *satellite*) may help to achieve the communicative goal of another one (the *nucleus*). This characterization comes close to the notion of *presentational relations* in RST (Mann & Thompson 1988: 257), whose terminology we borrow. Relations that belong to this category are defined in terms of the kind of goal the nucleus has, the way in which the nucleus potentially fails in achieving that goal by itself, and the way in which the satellite helps to solve that problem. For instance, if the nucleus is a directive (a request, a command, a piece of advice, an offer, etc.), whose goal is to get the addressee to perform a certain action, he/she may be reluctant or unable to perform it. *Motivation* (22) deals with reluctance; *Enablement* (23) deals with potential objective obstacles.

(22) You should eat more vegetables. Vegetables are good for you.

(23) Please can you post these letters? The stamps are in the drawer.

The goal of an assertion is to get the addressee to believe a proposition or accept it to a degree sufficient for the purposes of the exchange (Stalnaker 2002). If the addressee is reluctant to believe or accept the proposition, *Evidence* (24) can be provided to make it more plausible (adapted from Blakemore & Carston 2005).

(24) He must have been here recently. There are his footprints.

Justify is a relation, where the satellite is supposed to increase the addressee's readiness to accept the speaker's right to present the nucleus. Example (21) above could be considered an instance of *Justify*. Without the explanation in the second sentence the question might be considered a rude interference with the addressee's personal affairs.

Finally, in a *Background* relation (25), the satellite (ψ) is supposed to increase the addressee's ability to comprehend the nucleus (ϕ). [source of the example]

(25) [This piece begins with an anacrusis,] ϕ [an unaccented note which is not part of the first full bar.] ψ

Obviously, these relations can only connect speech acts, since their contribution to the coherence of the discourse is based crucially on the properties of speech acts. Therefore they cannot be embedded under operators like *if*. Indeed, there is no point in providing evidence for something you do not assert, or motivating a course of action you do not ask the addressee to commit to.⁴

In order to see how all this helps us to explain the scopal behaviour of *Elaboration* and *Explanation*, it is necessary to clarify some important aspects of the notion of source of coherence as we understand it. Notice that the definitions of semantic relations in (8) and those of presentational relations are mutually compatible. In fact, in most instances of presentational relations there will also be a causal or resemblance relation at some level. For example, (25) is an instance of *Background*, but the relation between the expressions *anacrusis* and *an unaccented note which is not part of the first full bar* can also be characterized as *Entity-Elaboration*. That is, the speaker uses *Elaboration* to provide *Background*. *Motivation*, *Evidence* and *Justify* almost always involve causality. For instance, in (22) the addressee should eat vegetables *because* they are healthy, and in (24), the speaker believes that ‘he’ must have been here recently *because* there are his footprints.

The reverse is not true: The coherence of a sequence can be due to semantic relations between the described situations, without there being any presentational component. In such cases, the speaker is simply describing a complex state of affairs that cannot be put into just one clause. The speaker is forced to split the content into manageable pieces and to order those pieces, and she uses the semantic categories behind (8) as guiding principles for splitting and linearization. The source of coherence of a given pair of discourse segments is the *speaker’s* reason to produce these segments next to each other in the given order. If the speaker is just describing different aspects of a complex state of affairs then the source of coherence is semantic. If the speaker is using one segment to help achieve the communicative goal of another then the source of coherence is pragmatic, presentational, or speech-act-level, whichever you like to call it.

Looking back at the relations in (8) as guidelines for splitting and ordering complex content, it is not difficult to notice that *Elaboration* and *Explanation* stick out as rather counterintuitive. *Elaboration* is a relation between descriptions of the same situation (or entity, in the case of *Entity-Elaboration*). But what is the point

⁴ The only apparent exception from this generalization is *Background*. Even if the phrase *an unaccented note which is not part of the first full bar* in (25) helps understanding the entire preceding sentence, it really only targets the expression *anacrusis*, which needs clarification. This expression can be embedded under *if*, as in *If this piece began with an anacrusis, an unaccented note which is not part of the first full bar, ...* However, the clarification only works if the speaker commits to the proposition that an *anacrusis* is an unaccented note which is not part of the first full bar. This commitment is communicated by an independent speech act. In other words, the satellite of a *Background* relation must be a speech act in any case.

of describing the same thing twice? As convincingly argued by Blakemore (1993, 1997), there must always be some other reason to produce an *Elaboration*, such as *Background* in (25) which helps understand the first utterance.

As for *Explanation*, i.e. the backwards causal relation, the question arises of why anyone who is simply trying to present a causally connected chain of events would do it backwards, while they have *Narration* and *Result* as more intuitive discourse strategies at their disposal? Of course, the speaker could forget to mention some previous event, realize later that without it the story does not make sense or sounds unconvincing, and try to repair it. But then, what the speaker is effectively doing is providing *Background* or *Evidence*. Our hypothesis is that backwards causal relations are generally used for presentational purposes.

In other words, an NRC that appears to elaborate on the host clause or to stand in a backwards causal relation to it (as in (1) discussed in section 4.1, cf. (16), (18) and (19)), will most of the time also provide *Background*, *Evidence* or some other kind of presentational support for some speech act in the context. That means that its real attachment site in the discourse structure is not local, that it represents a speech act itself and therefore cannot appear in the scope of entailment-canceling operators.

4.3 Conjunction and juxtaposition

It turns out that *Elaboration* and *Explanation* differ from *Narration*, *Result*, *Parallel* and *Contrast* also in another respect. Relations of the latter group are most naturally expressed with a conjunction (*but* in the case of *Contrast*). Removing the conjunction and replacing it by a period may slightly reduce the naturalness of the discourses, but this is easily repaired by adding a discourse adverbial that explicitly signals the relation, such as *then* for *Narration* (26b), or *also* for *Parallel* (27b). However, whether with or without a conjunction, the inferred coherence relation remains intact.

- (26) a. Peter called the dean, *and* she called the chair. *Narration*
b. Peter called the dean. (Then) She called the chair.
- (27) a. Mary should stand in front of Peter. *Parallel*
and he should move a little closer to Max.
b. Mary should stand in front of Peter.
(Also,) He should move a little closer to Max.

Elaboration and *Explanation* show a different pattern: The most natural way of expression is plain juxtaposition of independent sentences without any explicit markers. Adding *and* may change the interpretation dramatically. In (28b), from Carston 1993, the sentences are not understood to be about the same actress anymore, and in (29b) the temporal order is reversed as compared to (29a) (Asher & Vieu 2005).

- (28) a. I met a great actress at the party. I met Vanessa Redgrave. *Elaboration*
 b. I met a great actress at the party. *And* I met Vanessa Redgrave. *Parallel*
- (29) a. John fell. Chris pushed him. *Explanation*
 b. John fell, *and* Chris pushed him. *Narration*

This pattern is usually attributed to the semantics of the natural language conjunction *and*, which is not equivalent to logical conjunction (Carston 1993; Blakemore & Carston 1999, 2005; Txurruka 2003; Zeevat & Jasinskaja 2007). In particular, Txurruka (2003) argues that the natural language conjunction *and* is only compatible with coordinating relations in SDRT sense, which include *Narration*, *Parallel*, etc. In contrast, *Elaboration* and *Explanation* are most typical representatives of subordinating relations. (More on this distinction in section 5.) Zeevat & Jasinskaja (2007) propose that *and* requires the propositions it connects to serve as distinct answers to the same topic question, such as *Who did the speaker meet at the party?* in (28b), *What happened to John?* in (29b). Alternatively, that topic question can be seen as an abstract of the content of the conjoined discourse segments, implemented in the notion of Coordinated Discourse Topic (CDT) in SDRT (Txurruka 2003). All coordinating coherence relations must have a CDT, and the inference of such relations requires finding an appropriate CDT, or topic question. That means that inferring a coordinating relation in the absence of an explicit *and* basically involves ‘inferring an implicit *and*’.

In section 3 we came forward with the thesis that, unlike coordinate clauses connected by *and* or *but*, NRCs do not give you a conjunction ‘for free’. While *and* and *but* tell us explicitly how to combine the propositions expressed by their conjuncts in the semantics, NRCs do not generally tell us how to combine the propositions of the host clause and the NRC. The case is delegated to pragmatics, and only if we can infer a coherence relation between the NRC and the host clause that conjoins them at propositional level, does the NRC contribute locally to the semantics of the sentence. Coordinating coherence relations contain an equivalent of natural language conjunction as part of their semantics. This is why NRCs can be interpreted locally, as in (2) and (7), if they are attached to the host clause by *Narration*, *Parallel*, *Contrast*, etc. Since coordination with *and* or *but* still remains a preferred way of expression for these relations, local readings of the respective NRCs remain less salient and probably slightly degraded (about 50% availability in Poschmann’s 2018 experiments). But just like in (26), adding an adverbial improves their acceptability (up to 83% in Poschmann’s 2018 experiments).

Elaboration and *Explanation* work differently. Blakemore & Carston (1999, 2005) have argued that the purely semantic essence of these relations—identity or part-whole relations between the described situations in the case of *Elaboration* and backwards causality in the case of *Explanation*—are in principle compatible with

and. In a carefully constructed context like (30), there may be a backwards causal relation between the conjuncts of *and*, but it is clear that the cause (John tripped on a Persian rug) is not used to support, i.e. provide *Evidence*, *Background* for, or *Justify* the effect (John broke his leg). Rather, the conjoined clauses stand on a par as arguments against A's claim that the mats are not dangerous. Since the backwards causal relation is not accompanied by a presentational one, nothing prevents it from holding in the scope of *if*, with *and* (31), or even with an NRC (32).

- (30) A: Bob wants me to get rid of these mats. He says he trips over them all the time. Still, I don't suppose he'll break his neck.
B: Well, I don't know.
JOHN | broke his leg | and HE | tripped on a PERSIAN RUG |
- (31) If Bob breaks his leg and he trips on one of your mats, you will have to pay his medical bills.
- (32) If this happens to Bob, who trips on one of your mats, you will have to pay his medical bills.

Presumably, this is what is going on in exceptions like (33) where an NRC appears to be attached locally by *Elaboration* (note that (33) can also be paraphrased with *and* without change in interpretation).

- (33) If you get an email from Bill, who writes that he got a job, don't trust it.

This is an instance of purely semantic *Elaboration*, that is only concerned with mereological relations between events and does not carry any additional presentational functions. However, cases like (33) and (32) are exceptions, rather than the rule. As we argued in section 4.2, *Elaboration* and *Explanation* normally go together with a presentational relation and therefore connect speech acts. Presumably, this is also one of the reasons why they are so naturally expressed by juxtaposed independent sentences. In sum, it is not the semantic content of *Elaboration* and *Explanation* that is incompatible with natural language conjunction *and* and unsuitable for local attachment of NRCs, but their typical pragmatic function as *Evidence*, *Motivation*, *Justify* and *Background*.

5 Why do NRCs project most of the time?

In the previous section, we explained why NRCs connected to their host clause by a coordinating coherence relation (*Result*, *Narration*, *Parallel*, *Contrast*) are interpreted locally, and (normally) interpreted globally if they express a so-called subordinating relation, like *Elaboration* or *Explanation*. But why is the global interpretation of NRCs so pervasive, that until very recently we even believed that

they always project? Our answer is: Because NRCs are syntactically subordinate clauses (an a par with adverbial and complement clauses), and subordinate clauses express subordinating coherence relations most of the time.

The distinction between subordinating and coordinating relations in discourse structure should not be confused with the syntactic notion of subordination (Blühdorn 2008). Subordinate clauses are distinguished from matrix clauses in terms of specific syntactic behavior and formal properties (e.g. final position of the verb in German). Subordinating coherence relations differ from the coordinating ones not only in that they do not have CDTs, but also in the way they influence the prominence and accessibility of attachment points in discourse structure and antecedents for anaphora. Roughly, discourse-structurally subordinate material is relatively less prominent. In general, the interpretation of discourse proceeds from left to right, and the last-processed, rightmost, discourse unit is most prominent and most accessible for the attachment of new discourse material. However, if the most recent discourse unit is connected to the previous one by a subordinating relation, then the less recent superordinate unit is relatively more accessible (the generalization referred to in the literature as the Right Frontier Constraint, going back to Polanyi 1988; Webber 1991). That is, it can still serve as a discourse attachment point and provide antecedents for anaphora. In (34) for example, the first two sentences are related by *Parallel*, which is coordinating. Since the discourse proceeds from left to right, the most accessible referent for the pronoun *she* is the subject of the sentence preceding the pronoun, *Sue*. In (35), by contrast, the first two sentences are related by *Explanation*, which makes the second sentence subordinate to the first. Accordingly, both *Mary* and *Sue* are possible antecedents for the pronoun *she*.

- (34) Mary repaired the roof. Sue painted the wall. She_{Sue} ... *Parallel*
 (35) Mary repaired the roof. Sue told me so. She_{Mary/Sue} ... *Explanation*

Syntactic subordination has a similar effect on prominence. A pronoun is less likely to refer to an antecedent in a subordinate clause, than to one in a main clause, even if the subordinate clause is more recent (see Clark & Sengul 1979: exp. 3, Cooreman & Sanford 1996, and Frazier & Clifton 2005 for related experimental studies):

- (36) a. Mary left before Sue arrived. She_{Mary > Sue} ... main-subordinate
 b. Mary left. Then Sue arrived. She_{Sue} ... main-main
 c. After Mary left, Sue arrived. She_{Sue} ... subordinate-main

Syntactic and discourse subordination need not go together. Discourse subordination is possible without syntactic subordination, as in (35), as well as syntactically subordinate clauses need not necessarily express subordinating coherence relations. However, it seems that subordinate clauses are particularly well suited to encode

discourse-structural subordination. This is manifested in a statistical tendency for subordinate clauses to realize subordinating coherence relations (Matthiesen & Thompson 1988: 308), and non-restrictive relative clauses are no exception in this respect. In Loock's (2007) study of 450 English NRCs collected from texts of four different genres, only 20, i.e. 4.5%, were continuative, i.e. discourse-structurally coordinated (Holler 2008). Even if we do not yet understand the factors that determine whether a subordinate clause will be used to express a coordinating coherence relation, we have no reasons to believe that those factors will lead to an entirely different distribution under entailment-canceling operators.

In sum, NRCs express subordinating coherence relations like *Elaboration* and *Explanation* most of the time because their syntactically subordinate nature perfectly matches the prominence-related properties of these relations. As we have argued in the previous section, *Elaboration* and *Explanation*, in turn, tend to be used in connection with a presentational relation that requires that the NRC functions as an independent speech act and therefore cannot be embedded under operators like *if*.

6 Conclusions and outlook

In this paper, we started out with a set of previously unexplained observations about the scopal behavior of non-restrictive relative clauses, which turned out to allow for local interpretation in the scope of a conditional in an even greater variety of cases than previously believed. The theoretical explanation of NRC scope that we have developed is based on a minimal number of assumptions. Our central claim concerning the semantics of NRCs is that they (by themselves) do not provide any semantic link to the proposition expressed by the host clause, not even a plain conjunction. That link needs to be inferred pragmatically, in the same way as coherence relations are inferred between independent sentences in discourse. The rest of the story comes from the theory of discourse coherence. The way coherence inferences work in general and the properties of different classes of coherence relations explain why local interpretation is only possible for NRCs connected to their host clause by relations like *Narration* and *Parallel*, why relations like *Elaboration* and *Explanation* force global attachment and global interpretation of NRCs, and why NRCs in general show such a strong tendency for global interpretation. This study is an example of how discourse theory can be used to solve a problem in sentence semantics.

The approach we have developed makes predictions that go beyond the small set of data that we were able to discuss on these pages. In particular, it describes the behavior of non-restrictive relatives under entailment-canceling operators of any kind, but we have only tested it on conditionals. Further research is needed to see whether NRCs show the same pattern in other non-veridical environments. We also

did not have a chance to pay proper attention to the aspectual properties of NRCs, whose influence on their scope seems to go beyond the effects of coherence relations. This, too, as well as a precise formal implementation of our theory remains a task for the future.

References

- AnderBois, Scott, Adrian Brasoveanu & Robert Henderson. 2011. Crossing the appositive/at-issue meaning boundary. In *Semantics and Linguistic Theory (SALT) 20*, 328–346. doi:10.3765/salt.v20i0.2551.
- Arnold, Doug. 2007. Non-restrictive relatives are not orphans. *Journal of Linguistics* 43. 271–309. doi:10.1017/S0022226707004586.
- Asher, Nicholas & Laure Vieu. 2005. Subordinating and coordinating discourse relations. *Lingua* 115. 591–610. doi:10.1016/j.lingua.2003.09.017.
- Blakemore, Diane. 1993. The relevance of reformulations. *Language and Literature* 2. 101–120. doi:10.1177/096394709300200202.
- Blakemore, Diane. 1997. Restatement and exemplification: A relevance-theoretic reassessment of elaboration. *Pragmatics and Cognition* 5. 1–19. doi:10.1075/pc.5.1.04bla.
- Blakemore, Diane & Robyn Carston. 1999. The pragmatics of *and*-conjunctions: The non-narrative cases. *UCL Working Papers in Linguistics* 11. 1–20.
- Blakemore, Diane & Robyn Carston. 2005. The pragmatics of sentential coordination with *and*. *Lingua* 115. 569–589. doi:10.1016/j.lingua.2003.09.016.
- Blühdorn, Hardarik. 2008. Subordination and coordination in syntax, semantics and discourse: Evidence from the study of connectives. In Cathrine Fabricius-Hansen & Wiebke Ramm (eds.), *'Subordination' versus 'Coordination' in Sentence and Text: A Cross-linguistic Perspective*, 59–85. John Benjamins. doi:10.1075/slcs.98.04blu.
- Carston, Robyn. 1993. Conjunction, explanation and relevance. *Lingua* 90. 27–48. doi:10.1016/0024-3841(93)90059-6.
- Chierchia, Gennaro. 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In Adriana Belletti (ed.), *Structures and Beyond, Volume 3: The Cartography of Syntactic Structures*, 39–103. OUP USA.
- Clark, Herbert H. & C. J. Sengul. 1979. In search of referents for nouns and pronouns. *Memory & Cognition* 7(1). 35–41. doi:10.3758/BF03196932.
- Cooreman, Anne & Anthony J. Sanford. 1996. Focus and syntactic subordination in discourse. Research paper 79, HCRC, Edinburgh, Scotland, U.K.
- Del Gobbo, Francesca. 2003. *Appositives at the interface*: University of California, Irvine PhD dissertation.

- Espinal, M. Teresa. 1991. The representation of disjunct constituents. *Language* 67. 726–762. doi:10.2307/415075.
- Fabb, Nigel. 1990. The difference between English restrictive and non-restrictive relative clauses. *Journal of Linguistics* 26. 57–78. doi:10.1017/S0022226700014420.
- Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998–998. doi:10.1126/science.1218633.
- Frazier, Lyn & Charles Clifton. 2005. The syntax-discourse divide: Processing ellipsis. *Syntax* 8(2). 121–174. doi:10.1111/j.1467-9612.2005.00077.x.
- Geurts, Bart. 2010. *Quantity Implicatures*. Cambridge University Press.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell.
- Hobbs, Jerry R., Mark E. Stickel, Douglas E. Appelt & Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence* 63(1). 69–142. doi:10.1016/0004-3702(93)90015-4.
- Holler, Anke. 2008. A discourse-relational approach to continuation. In Anton Benz & Peter Kühnlein (eds.), *Constraints in Discourse*, 249–266. John Benjamins. doi:10.1075/pbns.172.12hol.
- Jasinskaja, Katja & Elena Karagjosova. forthcoming. Rhetorical relations. To appear in: Matthewson, Lisa & Cécile Meier & Hotze Rullmann & Thomas Ede Zimmermann (eds.): *The Companion to Semantics*. Oxford: Wiley.
- Kehler, Andrew. 2002. *Coherence, Reference, and the Theory of Grammar*. CSLI Publications.
- Koev, Todor K. 2013. *Apposition and the structure of discourse*: Rutgers University-Graduate School-New Brunswick PhD dissertation.
- Loock, Rudy. 2007. Appositive relative clauses and their functions in discourse. *Journal of Pragmatics* 39(2). 336–362. doi:10.1016/j.pragma.2006.02.007.
- Mann, William C. & Sandra Thompson. 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text* 8(3). 243–281. doi:10.1515/text.1.1988.8.3.243.
- Martin, Scott. 2016. Supplemental update. *Semantics and Pragmatics* 9. doi:10.3765/sp.9.5.
- Matthiesen, Christian & Sandra A. Thompson. 1988. The structure of discourse and subordination. In John I. Haiman & Sandra A. Thompson (eds.), *Clause Combining in Discourse and Grammar*, 275–329. Amsterdam: Benjamins. doi:10.1075/tsl.18.12mat.
- McCawley, James D. 1982. Parentheticals and discontinuous constituent structure. *Linguistic Inquiry* 13. 91–106.
- Nouwen, Rick. 2007. On appositives and dynamic binding. *Journal of Language and Computation* 5(1). 87–102. doi:10.1007/s11168-006-9019-6.

- Polanyi, Livia. 1988. A formal model of the structure of discourse. *Journal of Pragmatics* 12. 601–638. doi:10.1016/0378-2166(88)90050-1.
- Poschmann, Claudia. 2018. Embedding non-restrictive relative clauses. In Uli Sauerland & Stephanie Solt (eds.), *Sinn und Bedeutung* 22, vol. 2 ZASPiL 61, 235–252. Berlin: ZAS.
- Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford University Press.
- Sæbø, Kjell Johan. 2011. Appositives in modal contexts. In Ingo Reich et al. (ed.), *Sinn und Bedeutung* 15, 79–100. Saarland University Press.
- Safir, Ken. 1986. Relative clauses in a theory of binding and levels. *Linguistic Inquiry* 17(4). 663–689.
- Sanders, Ted J. M., Wilbert P. M. Spooren & Leo G. M. Noordman. 1992. Toward a taxonomy of coherence relations. *Discourse Processes* 15(1). 1–35. doi:10.1080/01638539209544800.
- van der Sandt, Rob A. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9(4). 333–377. doi:10.1093/jos/9.4.333.
- Schlenker, Philippe. 2013. Supplements without bidimensionalism. Unpublished manuscript. Institut Jean-Nicod and NYU. https://www.semanticsarchive.net/Archive/jgwMjNmM/Supplements_without_Bidimensionalism.pdf.
- Simons, Mandy, Judith Tonhauser, David Beaver & Craige Roberts. 2011. What projects and why. In *Semantics and Linguistic Theory (SALT) 20*, 309–327. doi:10.3765/salt.v20i0.2584.
- Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25(5). 701–721. doi:10.1023/A:1020867916902.
- Sweetser, Eve. 1990. *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge University Press.
- Txurruka, Isabel Gómez. 2003. The natural language conjunction *and*. *Linguistics and Philosophy* 26. 255–285. doi:10.1023/A:1024117423963.
- Webber, Bonnie Lynn. 1991. Structure and ostension in the interpretation of discourse deixis. *Natural Language and Cognitive Processes* 2(6). 107–135. doi:10.1080/01690969108406940.
- Zeevat, Henk. 2016. Local satisfaction explained away. In *Semantics and Linguistic Theory (SALT) 26*, 264–283. doi:10.3765/salt.v26i0.3790.
- Zeevat, Henk & Katja Jasinskaja. 2007. *And* as an additive particle. In M. Aurnague, Kepa Korta & Jesus M. Larrazabal (eds.), *Language, Representation and Reasoning. Memorial volume to Isabel Gómez Txurruka*, 315–340. University of the Basque Country Press.

Attachment in Syntax and Discourse: Scope of NRCs

Katja Jasinskaja
Universität zu Köln
SFB 1252 *Prominence in Language*
Luxemburger Str. 299
50939 Cologne
Germany
katja.jasinskaja@uni-koeln.de

Claudia Poschmann
Goethe-Universität Frankfurt
Institut für Linguistik
Norbert-Wollheim-Platz 1
60629 Frankfurt/M.
Germany
claudia.poschmann@gmx.net