Questions and Goals in the Structure of Discourse: A new look at coherence relations

Dr. Katja Jasinskaja

geb. am 1.7.1976
in Moskau, Russland

Kumulative Habilitationsschrift
eingereicht bei der Philosophischen Fakultät der Universität zu Köln
Mai 2017

Contents

1 Questions and goals in the structure of discourse **5**
1.1 Background 5
1.2 Exhaustivity and topic continuity: *Elaboration* as default 16
1.3 Marking coordinating relations 20
1.4 Correction 25
1.5 Grounding and acceptance: Explaining subordinating relations 29
1.6 Further applications of the QUD approach 32
1.7 Conclusion 39

2 Modelling discourse relations by topics and implicatures **47**
Jasinskaja, E., 2010b. Modelling discourse relations by topics and implicatures: The elaboration default. In: Benz, A., Kühnlein, P., Sidner, C. (eds.), 2010. Constraints in Discourse 2. John Benjamins, Amsterdam, pp. 61–79.

3 Explaining conjunction systems: Russian, English, German **69**
Jasinskaja, K., Zeevat, H., 2009. Explaining conjunction systems: Russian, English, German. In: Riester, A., Solstad, T. (eds.), *Proceedings of Sinn und Bedeutung 13*. Vol. 5 of SinSpeC, pp. 231–245.

4 Explaining additive, adversative, and contrast marking in Russian and English **87**
Jasinskaja, K., Zeevat, H., 2008. Explaining additive, adversative and contrast marking in Russian and English. *Revue de Sémantique et Pragmatique* 24, 65–91.

5 Corrective constrast in Russian, in contrast **113**
Jasinskaja, K., 2010a. Corrective contrast in Russian, in contrast. *Oslo Studies in Language* 2 (2), 433–466.

6 Correction by adversative and additive markers **147**
Jasinskaja, K., 2012. Correction by adversative and additive markers. *Lingua* 122, 1899–1918.

7 Corrective elaboration **183**
Jasinskaja, K., 2013. Corrective elaboration. *Lingua 132*, 51–66.

8 Rhetorical relations **211**
Jasinskaja, K., Karagjosova, E., submitted. Rhetorical relations. Submitted to: Matthewson, L., C. Meier, H. Rullmann, T. E. Zimmermann (eds.): *The Companion to Semantics*. Oxford: Wiley.

9 Discourse-level implicature: A case for QUD **241**
Jasinskaja, K., Salfner, F., Freitag, C., 2017. Discourse-level implicature: A case for QUD. *Discourse Processes* 54 (3), 239–258.

10 Not at issue any more **269**
Jasinskaja, K., in revision. Not at issue any more. Submitted to *Journal of Semantics*.

Chapter 1

Questions and goals in the structure of discourse

1.1 Background

Relational vs. goal-based approaches to discourse structure

What makes a text a text? What enables us to recognize that (1) is a text, whereas (2) is a random sequence of sentences?

(1) a. My cell phone is acting up.
b. I keep pressing the home button
c. but when I look around, I'm still at work.

(2) a. My cell phone is acting up.
b. Round is a shape
c. but I don't give out free samples.[1]

Inquiry into the nature of text has singled out a number of constitutive properties that distinguish texts from non-texts. These include cohesion, coherence, intentionality, acceptability, informativity and situationality, among others. The focus of the present volume is on the relationship between coherence and intentionality. *Coherence* refers to the property of internal connectivity of a text, or more generally, discourse. A discourse is coherent if meaningful links exist between all its parts (sentences, clauses). *Intentionality* concerns the relationship between a discourse and its goal. Here it is essential that a discourse realizes a sensible strategy in achieving its goal. This volume addresses the question which property of discourse is more important from a linguistic point of view. Is it coherence or intentionality of the surrounding context that helps

[1] Both (1) and (2) are taken from an online collection of sayings at http://coolfunnyquotes.com (last accessed on May 6, 2017) but if (1) is an intact saying, (2) is picked out clause by clause from three different sayings.

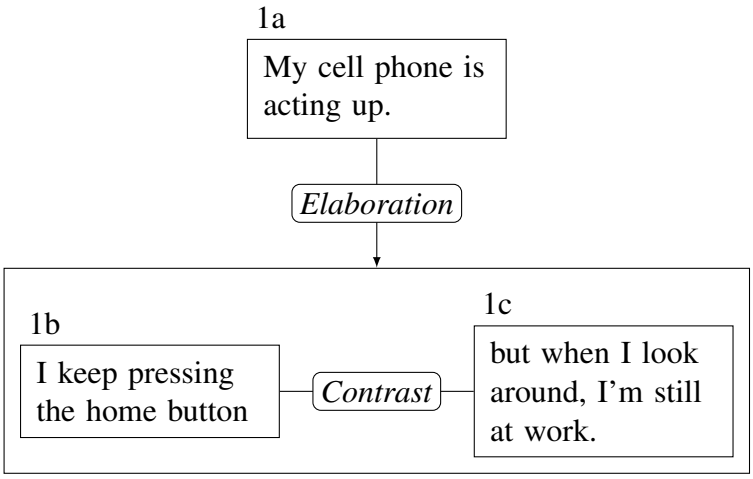

Figure 1.1: Relational structure for (1)

us better understand the functioning of context-sensitive linguistic expressions and grammatical constructions? Or is it even possible to characterize coherence without reference to intentionality?

Research on discourse structure in linguistics gave rise to two kinds of positions on this matter. According to one position, let's call it the *relational* approach, represented by the work of e.g. Hobbs (1985), Mann and Thompson (1988), Kehler (2002) and Asher and Lascarides (2003), linguistically most relevant property of discourse is coherence, and to describe the structure of a discourse means primarily to give a characterization of the 'meaningful links' between its parts. These links are referred to by different authors as *coherence relations*, *rhetorical relations*, or *discourse relations*.

For instance, (1-b) and (1-c) taken together describe the same problematic situation as the first sentence in (1-a) does, but they describe it in more detail. Therefore (1-b)–(1-c) is an *Elaboration* of (1-a). The connective *but* in (1-c) indicates that this clause denies some implication of (1-b), presumably the ungrounded expectation of the speaker that pressing the home button would physically take him home. The 'meaningful link' between (1-b) and (1-c) is *Contrast*—a relation between a claim that goes in favour and a claim that goes against some implicit proposition. The resulting relational structure is shown in Figure 1.1.

Elaboration and *Contrast* are coherence relations. Other coherence relations are *Explanation* (one discourse unit gives the cause of the situation described in the other), *Narration* (relation between descriptions of events occurring in a sequence in narrative discourse), *Parallel* (the similarities between two situations are emphasized), and a number of others. An overview of commonly assumed types of coherence relations is given in chapter 8 (Jasinskaja and Karagjosova, submitted). As for (2), this sequence of sentences is incoherent because we are not able to recognize that the clauses could stand in any of these relations to each other.

Another approach to discourse structure is what I will refer to as the *goal-based* approach (developed, for instance, by Grosz and Sidner, 1986). It puts emphasis on the requirement that a discourse should pursue some recognizable communicative goal and one should be able to conceive of each part of the discourse as a reasonable step in achieving the overall goal. In other words, as pointed out above in relation to the notion of intentionality, the discourse must realize a sensible *strategy* to achieve that goal.

Assuming that the speaker of (1) is sincere and is not trying to make a joke, his ultimate (domain-level) goal is probably to get help with his cell phone problem. That requires informing the hearer about the problem. This is the overall *communicative goal* of (1). The first sentence (1-a) clearly pursues that goal, but it does not give enough information for anyone to be able to help. The goal of (1-b) and (1-c) is to fill this gap. These clauses achieve it by describing the inconsistent behaviour of the cell phone with respect to its supposed function of getting the speaker home. The communicative goal of (1-b) is to describe the speakers actions directed at getting home, whereas the goal of (1-c) is to describe a situation that is opposite to being home, which in sum gives rise to the perceived inconsistency. Thus (1) makes sense because we can figure out what the speaker is trying to achieve and how each of the things he says contributes to achieving that goal.

In contrast, it is not clear why someone would be saying (2). And even if we know from the situational context that the speaker is trying to get help with his phone, it is not clear how (1-b) and (1-c) contribute to that goal. That is why (2) does not make sense.

The papers in this collection contribute to a research programme that tries to find answers to the questions: Which approach to discourse structure is most useful from a linguistic point of view? Should we characterize discourse context in terms of coherence relations or in terms of communicative goals in order to better describe the properties of context sensitive linguistic devices? Do we, perhaps, need both kinds of structure, or can relational structure be derived from goal structure or vice versa?

Formal methods in sentence and discourse semantics

The choice of theoretical methodology for the present research programme is dictated by the commitment to precision and predictive power. The last fifty years have seen the rise of formal methods such as mathematical logic and model theory in sentence semantics, following the seminal work of Richard Montague (1970, 1973). Precise semantic theories have been developed for tense and aspect (Dowty, 1979), quantification (Barwise and Cooper, 1981), plurals (Link, 1983), modality (Kratzer, 1981), conjunction (Partee and Rooth, 1983), to name just a few classical examples. However, natural languages possess a whole range of expressions whose interpretation is dependent on context. These include anaphoric pronouns (*he*) and adverbials (*then*), tense, ellipsis, intonation, as well as any linguistic expressions that carry presupposi-

tions, such as phase verbs (*stop*) and factives (*regret*). Moreover, phenomena that go beyond the conventional meaning of linguistic expressions, i.e. pragmatic inferences such as conversational implicatures (*some* $\rightsquigarrow$ *not all*), are intrinsically context sensitive. Models of sentence semantics fall short when it comes to describing the semantics of context-sensitive expressions. In cases where context intrudes into the meaning of a sentence the typical way out is to assume that the semantic representation of the sentence contains a variable whose value is provided by the context. However, what kind of value is to be expected in what kind of context generally remains an issue outside the scope of such theories.

One of the goals of discourse semantics is to provide a theory of discourse context that can fill the holes in sentence meaning. In order to be able to communicate with existing formal semantic descriptions, a theory of discourse must be formulated at a comparable level of precision. Substantial progress has been made in the development of such a theory implementing in particular the ideas of the relational approach. Most notably, Segmented Discourse Representation Theory (SDRT, Asher and Lascarides, 2003) provides a rich formalism that allows to construct semantic representations of discourses, including coherence relations that hold between sentences. This approach has been successfully applied to a whole range of context-sensitive linguistic phenomena including tense (Lascarides and Asher, 1993), lexical disambiguation (Asher and Lascarides, 1995), ellipsis (Asher et al., 1997), definites (Asher and Lascarides, 1998a), presupposition (Asher and Lascarides, 1998b), embedded implicatures (Asher, 2013), demonstrating how discourse structure based on coherence relations sets the relevant contextual parameters.

Another goal of discourse semantics is to describe the extra content that comes on top of the content of individual sentences once they are put together to form a coherent whole. Ideally, discourse semantics should also provide a theory of how that extra content is inferred based on linguistic and situational input. For instance in (1), the extra content includes the information that the problem with the cell phone is that it doesn't take the speaker home (subsumed under the label of *Elaboration*) and the speaker's expectation that pressing the home button would take him home (a consequence of *Contrast*). One of the main accomplishments of SDRT is that it provides a non-monotonic logic that models the inference of coherence relations from a number of knowledge sources including various aspects of sentence semantics, explicit cues such as *but*, as well as word knowledge. Combined with precise semantic definitions of coherence relations, that takes us all the way from linguistic and situational input to the 'extra content'. A record not easy to beat.

A formal goal-based theory?

Until now there has been no equally worked out and comprehensive formal theory of discourse based on communicative goals. There are both good reasons and challenging obstacles for developing such a theory. The good reasons have to do primarily with the

need to overcome some fundamental limitations of the relational approach.

The first problem of the relational approach is the need to justify the inventory of coherence relations. How many relations are there? Why these relations and not others? The lack of convincing answers to these questions has been the source of dissatisfaction with the approach ever since it appeared on the market, and a lot of effort has been invested in attempts to break down coherence relations into combinations of more basic independently motivated features (see Knott and Dale, 1994; Knott, 1996; Knott and Sanders, 1998). An overview of this discussion is given in section 2.2 of chapter 8 (Jasinskaja and Karagjosova, submitted). In sum, coherence relations are too complex to serve as ontological primitives of a theory, as this is the case, for instance, in SDRT. Whether a goal-based theory would do better on this point depends on the specific way it is implemented. It does not make much sense to develop a classification of communicative goals that would simply double the familiar types of coherence relations like *Elaboration*, *Explanation*, etc. (cf. discussion in Hunter and Abrusán, forthcoming). The challenge is to develop a goal-based theory that runs entirely on independently motivated categories and principles.

The second limitation of the relational approach has attracted far less attention so far, but is becoming particularly relevant in the context of more recent research on discourse expectations (Kehler et al., 2008; Bott and Solstad, 2014). To illustrate this point, Benz and Jasinskaja (2017) cite the definition of discourse coherence formulated by Bill Mann, perhaps ironically, one of the first activists of the relational approach to discourse structure:

> ... For every part of a coherent text, there is some function, some plausible reason for its presence, evident to readers, and furthermore, there is no sense that some parts are somehow missing. (Mann, the RST website)[2]

Somewhat reformulating Mann's subsequent remarks about Rhetorical Structure Theory (RST, Mann and Thompson, 1988), one could say that rhetorical relations capture the first aspect of discourse coherence—an evident role for every part. But what about the second aspect—no parts missing? Strikingly, all theorizing about discourse has concentrated on the first aspect of discourse coherence so far. Existing formal definitions of coherence (Asher and Lascarides, 2003; Ginzburg, 2012) measure it by the availability and the strength of meaningful links between utterances *as they stand*. None of them captures the intuition that the discourse in (3) as it stands, without a continuation like that in (4), sounds odd. For instance, an SDRT analysis would recognize an *Elaboration* relation between (3-b) and (3-a). Since the sentences are rhetorically connected, the discourse (3) is coherent, and just as coherent as (3) followed by (4), which constitutes another *Elaboration* of (3-a). It is hard to see how a notion of coherence based exclusively on coherence relations would capture the fact that the information about the second "thing" done by the speaker on his birthday is missing in (3).

[2]http://www.sfu.ca/rst/01intro/intro.html, last accessed on May 10, 2017.

(3) John Scalzi, *Old Man's War*:

a. I did two things on my seventy-fifth birthday.

b. I visited my wife's grave.

(4) Then I joined the army.

In order to be able to see that something is missing, one has to have a notion of what is needed, i.e. of the *goal* the discourse is supposed to fulfil.

There is a growing body of research showing that different linguistic devices generate different kinds of expectations with respect to the upcoming discourse, most notably the work on the implicit causality of verbs (Kehler et al., 2008; Bott and Solstad, 2014), showing that causality in the semantics of a verb makes readers/hearers expect an upcoming *Explanation*. While it is certainly possible to enrich a theory based on coherence relations with a special forward-looking mechanism,[3] "expectations" do not naturally fit into a standard relational framework. In contrast, communicative goals provide a source for missing, expected, or needed information in the most straightforward way.

The main obstacle for the development of a goal-based theory of discourse structure at a level of precision comparable to SDRT has long been the absence of a formal semantics for goals. The early goal-based proposals (Grosz and Sidner, 1986; Litman and Allen, 1990) did not aim for a fully formalized model of truth-conditional discourse semantics and did not see a worked out interface to compositional sentence semantics as their highest priority (see Asher and Lascarides, 2003, pp. 82–90, for detailed discussion). However, recent years have seen the rise of a group of approaches that promise a breakthrough in this area. These approaches replace the original goals by *Questions under Discussion* (QUD, e.g. Ginzburg, 1996), which turn out to be a more promising concept for the semanticist.

Questions under discussion as goals

The goal of an informative utterance can be seen as to answer a *Question under Discussion* (QUD, e.g. Klein and von Stutterheim, 1987; van Kuppevelt, 1995; Ginzburg, 1996; Roberts, 1996; Büring, 2003; Simons et al., 2011; Onea, 2016).[4] The question can be asked explicitly by another conversation participant, or it can remain implicit, in which case it must be understood from the context. The goal structure of (1) described

[3] SDRT uses underspecified variables as such a mechanism. For example, implicit causality verbs like *annoy* introduce a variable for the cause of annoyance in their lexical entry. That variable must be specified by context, which leads to the expectation of an *Explanation* relation (Danlos, 2001; Asher and Lascarides, 2003). The *Maximize Discourse Coherence* (MDC) principle establishes a preference for structures where all underspecified variables are resolved.

However, this works only as long as we can associate an expectation with a specific lexical item. In (3), for instance, it is less obvious which word could be responsible for the expected *Elaboration*.

[4] A vast bibliography on the concept of QUD and its various applications is maintained at http://www.ling.ohio-state.edu/~croberts/QUDbib/.

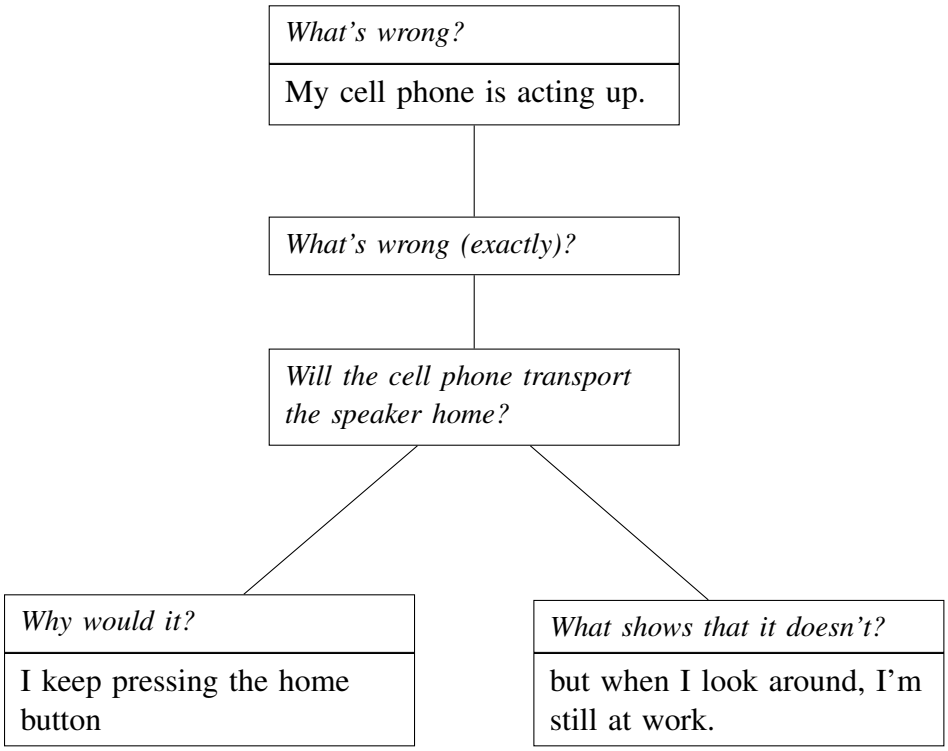

Figure 1.2: QUD structure for (1)

in the beginning of this section can then be expressed in terms of QUDs roughly as shown in figure 1.2. Even if the notion of QUD might not provide the most general characterization of communicative goals, it does provide a viable operationalization of that concept and, as Benz and Jasinskaja (2017) argue, is an attractive theoretical construct for a number of reasons.

First of all, the past decades have brought substantial progress in our understanding of the semantics of questions and the question-answer relationship (Hamblin, 1973; Karttunen, 1977; Groenendijk and Stokhof, 1984; Krifka, 2001; van Rooy, 2003; Groenendijk and Roelofsen, 2009). These studies provide a solid formal foundation for QUD-based theories. Relationships between goals can also be understood in terms of relationships between questions under this approach. For example, Grosz and Sidner (1986) define a *dominance* relation between goals:

> An action that satisfies one intention, say DSP1, may be intended to provide part of the satisfaction of another, say DSP2. When this is the case, we will say that DSP1 **contributes to** DSP2; conversely, we will say that DSP2 **dominates** DSP1 [...] (Grosz and Sidner, 1986, p. 179)[5]

[5]Grosz and Signer's 'DSP' is a *discourse segment purpose*. 'Intentions' and 'purposes' correspond to 'goals' in the present terminology.

But what does it mean to be 'part of the satisfaction of another intention'? Groenendijk and Stokhof (1984) define a notion of *question entailment* (5), cited from Groenendijk and Stokhof (2011, p. 1085), which serves as basis for Roberts' (1996) notion of *subquestion* (6):

(5) $?\phi$ entails $?\psi$ iff every possible answer to $?\phi$ entails some possible answer to $?\psi$.

(6) Q is a subquestion of Q' iff Q' entails Q.

The notion of dominance or subquestionhood captures an important aspect of what it means to have a sensible strategy in achieving a goal. An ambitious goal is easier to achieve if it can be divided into a number of subgoals each of which is relatively less ambitious. Similarly, a complex question can be answered by addressing a series of simpler subquestions, i.e. questions whose answers provide less information than the complex question requires, but which might provide a complete answer if the whole series of questions is resolved. While it remains open whether Groenendijk and Stokhof's and Roberts' subquestionhood captures all aspects of dominance originally intended by Grosz and Sidner (1986), it does provide a level of precision and is formulated in terms that formal theories of sentence semantics can interface with.

Second, questions reflect what is *relevant* at each point in discourse. They means that they stand in a systematic relation to the domain-level goals of the communication participants (van Rooy, 2003; Schulz and van Rooij, 2006). One and the same interrogative sentence can be intended to encode quite different questions, depending on the situation it is uttered in. This is illustrated by the example in (7) and (8), discussed by Ginzburg (1995) and van Rooy (2003). In (7) the question is really *In which city is Jill?*, because the destination city is what you want to know when you take a flight. In (8) the question is *At which address in Helsinki is Jill?*, because this is what you want to know when you take a taxi from the airport to your final destination.

(7) a. Context: Jill about to step off a plane in Helsinki.

Flight attendant: Do you know where you are?
Jill: Helsinki.

b. Flight attendant: Ah ok. Jill knows where she is.

(8) a. Context: Jill about to step out of a taxi in Helsinki.

Driver: Do you know where you are?
Jill: Helsinki.

b. Driver: Oh dear. Jill doesn't (really) know where she is.

Van Rooy (2003) provides a formal account of the semantics of questions which is dependent on the decision problem an agent is facing in the real world.

Finally, questions can also be seen as more or less linguistically defined "templates" for possible answers (as is the case, e.g., in Hamblin, 1973), which underlies

the notion of question-answer *congruence* and has proved useful in the analysis of a wide range of sentence-level phenomena. This includes first and foremost the information structure of the sentence, the accentuation pattern induced by the partition of the sentence into focus and background and the interpretation of focus-sensitive operators. The influential proposal of Roberts (1996) inspired much further theoretical work on information structure including Büring (2003) on contrastive topics and Beaver and Clark (2008) on focus particles. The intuition behind it is the same as that behind the well-known question-answer test used to detect the focus structure of a sentence. An answer to a question is appropriate only if its focused constituent corresponds to the *wh*-phrase of the question. For example, (10-a) with a nuclear pitch accent on the subject *Mary* (indicated by small caps) is an appropriate answer to (9-a), and not to the other questions in (9), and therefore the focus of (10-a) is on the subject *Mary*. Similarly, (10-b) only fits the question in (9-b) and instantiates focus on the direct object, whereas (10-c) is ambiguous between focus on *Sue* (9-c), focus on the whole VP (9-d), and broad focus on the whole sentence (9-e).

(9) a. Who introduced Bill to Sue?
b. Who did Mary introduce to Sue?
c. Who did Mary introduce Bill to?
d. What did Mary do?
e. What happened?

(10) a. MARY$_F$ introduced Bill to Sue.
b. Mary introduced BILL$_F$ to Sue.
c. Mary introduced Bill to [SUE]$_F$. /
Mary [introduced Bill to SUE]$_F$. / ...

Thus, a question determines an open proposition, e.g. $\lambda x[x$ introduced Bill to Sue$]$ in (9-a)/(10-a), and a *set of alternatives*, e.g. { 'Mary introduced Bill to Sue', 'John introduced Bill to Sue', 'Peter introduced Bill to Sue', ... }. Questions under discussion as open propositions or alternative sets serve as a bridge to the linguistic form of utterances and provide the kind of contextual input expected by many accounts of sentence semantics and sentence-level pragmatic inferences, as for instance in Rooth's (1992) account of focus semantics or in Kratzer and Schimoyama's (2002) and Chierchia's (2006) account of indefinites and free choice items.

One example of sentence-level pragmatic inferences that depend on contextually given alternative sets is discussed at some length in chapter 9 (Jasinskaja et al., 2017). It is Gricean quantity implicatures, in particular scalar and exhaustivity implicatures, that are inferred by reasoning with the Gricean Maxim of Quantity (11), restricted by Quality (12) and Relevance (13) (Grice, 1975):

(11) Maxim of Quantity:
a. Make your contribution as informative as required (for the current purposes of the exchange).

b. Do not make your contribution more informative than is required.

(12) Maxim of Quality:

a. Do not say what you believe to be false.
b. Do not say that for which you lack adequate evidence.

(13) Maxim of Relation:
Be relevant.

For example, the answer (10-b) to the question (9-b) has an exhaustivity implicature that Mary did not introduce anyone else beside Bill to Sue. Indeed, the speaker could have given a more informative answer, that Mary introduced Bill and John, or Bill and Peter to Sue, following the Quantity Maxim. The fact that the speaker did not give such an answer could mean either that information about John and Peter is not required for the current purposes of the exchange, i.e. irrelevant, or the speaker does not know for a fact that Peter or John were also introduced, so she does not mention them following the Quality Maxim. Assuming in addition that the speaker is competent, or opinionated on the question asked (otherwise why ask that person?), that means that the speaker in fact knows that John and Peter were not introduced.

For this inference it is crucial to have a notion of the other alternative things the speaker could have said. In this example, *Peter* and *John* are alternatives to *Bill*, and 'Mary introduced Bill and John to Sue' and 'Mary introduced Bill and Peter to Sue' are possible alternative utterances. The strength of the exhaustivity implicature depends, for instance, on whether the set of relevant alternatives only contains guests at a particular party, or all the people Mary has known in her entire life, and that, in turn, is restricted by the set of question alternatives. In other words, we need the QUD as input to calculate the exhaustivity implicature following the Gricean recipe.

To summarize, the notion of QUD provides an operationalization of the notion of communicative goal. It interfaces well with existing formal accounts of sentence semantics and sentence-level pragmatic inferences. Moreover, in combination with existing formal semantic theories of questions the approach makes it possible to develop a goal-based notion of discourse structure at a high level of precision.

However, until now there exists no systematic goal-based account of the full range of phenomena traditionally covered by relational approaches such as SDRT. These include, first and foremost, the inference of semantic effects of coherence relations—the 'extra content' hiding behind the labels of e.g. *Elaboration* and *Contrast* in example (1)/figure 1.1. Second, we need a goal-based theory that works for context sensitive linguistic phenomena, such as anaphora (*he, then*), presupposition (*stop, regret*), and implicature (*some* $\rightsquigarrow$ *all*), paying attention to the aspects of their behaviour that have been shown to be sensitive to coherence relations. In other words, if we want to show that discourse structure can be done in terms of communicative goals instead of coherence relations, we have to show that phenomena that have been explained in terms of coherence relations, can also be explained in terms of goals. Moreover, if we

do not want to reproduce the typical problems of the relational approach, that is, the need to justify a specific inventory of coherence relations, the goal-based theory we are looking for should not be based on a similar classification of goals, or QUDs. Instead, it should derive all the necessary categories from independently motivated categories and principles. This is the programmatic goal that has driven the work presented in this volume.

The contribution of the presented work

The present collection of papers demonstrates the development of an original new approach to discourse semantics that combines a goal-based, in particular QUD-based, notion of discourse structure with a handful of general independently motivated pragmatic principles to model the inference of the semantic and pragmatic effects of some major classes of coherence relations. In particular, it addresses two main questions:

1. Can the 'extra content' associated with common types of coherence relations in relational theories be inferred in a theory of discourse structure based on QUDs? Or put differently, can motivation for an inventory of coherence relations be drawn from QUD structure?

2. Does QUD-based discourse structure help explain context-sensitive linguistic phenomena in sentence semantics and pragmatics? How can it help resolve theoretical issues in the analysis of such phenomena that are the subject of current debate?

Chapters 2–8 concentrate primarily on the first question. They consider different classes of coherence relations, giving QUD-based accounts of pragmatic inferences or linguistic devices that establish those relations. It all starts with a pragmatic account of *Elaboration* as a default coherence relation in chapter 2 (Jasinskaja, 2010b), introduced in section 1.2 below. The next group of relations is *Contrast*, *Parallel* (or *List*), and *Narration* which I argue require at least some minimal linguistic signalling. Section 1.3 gives an introduction to the QUD-based semantic typology of discourse connectives that encode these relations developed in chapter 3 (Jasinskaja and Zeevat, 2009) and chapter 4 (Jasinskaja and Zeevat, 2008). The next coherence relation to be considered is *Correction*, which is a hybrid relation combining properties of *Elaboration* and *Contrast*. Section 1.4 gives an overview of three chapters, 5 (Jasinskaja, 2010a), 6 (Jasinskaja, 2012), and 7 (Jasinskaja, 2013), that elucidate different aspects of this relation. The mechanisms for the inference of *Elaboration* developed in chapter 2 are successfully applied to the inference of *Correction*. Finally, *Elaboration* is reconsidered in the context of other coherence relations, *Explanation* and *Background*, with which it shares (a) the ability to remain entirely implicit and (b) the characteristic effect it has on the prominence of antecedents for anaphora and accessibility of discourse units for attachment of new discourse material. Chapter 8 (Jasinskaja and

Karagjosova, submitted) extends the pragmatic framework of chapter 2 to provide a uniform account of the inference of these relations.

Some of these chapters also contribute to clarifying question 2. For instance, chapters 4 and 5 develop a QUD-based account of the characteristic information structure (focus, contrastive topic) and the typical accentual patterns of *Parallel*, *Contrast* and *Correction*. Chapter 8 sketches out an approach to modelling the influence of *Elaboration*, *Explanation* and *Background* on prominence and anaphora.

The last chapters 9 (Jasinskaja et al., 2017) and 10 (Jasinskaja, in revision), summarized in section 1.6, are devoted entirely to the second question. In particular, they concentrate on controversial issues in pragmatics concerning the nature of conversational and conventional implicatures, and demonstrate how the QUD-based view of discourse structure can shed new light on those issues. Chapter 9 argues that QUD-based discourse structure is needed for a Gricean account of exhaustivity implicatures above the sentence level. As explained in section 1.6, such implicatures indeed require a Gricean account and cannot be reduced to the effect of exhaustivity operators in the sentence structure *à la* Chierchia (2004). This has been a subject of recent fierce debate (cf. Geurts, 2010), also known as the localist/globalist debate. Chapter 10 demonstrates the effect of coherence relations on the pragmatic status of peripheral (not-at-issue) components of sentence content. It deals with the currently widely discussed issue of the variable at-issue status of non-restrictive (appositive) relative clauses and uses QUD-based discourse structure to bridge the gap between at-issue status and coherence relations.

1.2 Exhaustivity and topic continuity: *Elaboration* as default

Chapter 2 of this volume (Jasinskaja, 2010b) takes up the issue of how the semantic contribution of coherence relations is inferred in a framework without coherence relations. It concentrates in particular on *Elaboration* relations, illustrated in (14) and (15). The semantic contribution of *Elaboration* is that the states of affairs presented in the sentences connected by this relation are perceived as identical. That is, an *Elaboration* (in the broad sense assumed in the present body of work) describes the same eventuality in different words, in more detail, or from a different perspective. In (14), this leads in particular to the inference that the garment mentioned in the first sentence was a shirt, and it was damaged by staining. In (15), *Elaboration* implies that skis were Alena's main means of transport.

(14) Fred damaged a garment.
He stained a shirt.

(15) Alena broke her skis.
She lost her main means of transport.

The standard SDRT account of such inferences (Lascarides and Asher, 1993; Asher and Lascarides, 2003, pp. 204–207) relies on lexical and encyclopaedic knowledge about type–subtype and part–whole relations between semantic entities, e.g. *garment–shirt*, *damage–stain*. The presence of such relations triggers the application of a non-monotonic inference rule that establishes an *Elaboration* relation between the sentences. Once established, the semantic definition of *Elaboration* applies and imposes part–whole and identity relations between the eventualities and objects involved.

In contrast, the account of *Elaboration* developed in chapter 2 is based on two general pragmatic principles: the principle of exhaustivity (16), and the principle of topic continuity (17). Exhaustivity is a standard type of pragmatic enrichment of the literal meaning of a sentence (an implicature) which is sensitive to the current question under discussion. For example, as an exhaustive answer to the question *What did Alena break?*, the utterance *Alena broke her skis* implies that Alena broke her skis and nothing else. If the same utterance is interpreted exhaustively with respect to the question *What happened to Alena?*, then it implies that breaking her skis is the only (relevant) thing that happened to Alena.

(16) *The Principle of Exhaustivity:*
By default, an utterance is interpreted exhaustively (with respect to the current discourse topic).

(17) *The Principle of Topic Continuity:*
By default, the discourse topic does not change.

The second principle (17) states that by default, the discourse topic does not change, i.e. in the absence of explicit indications to the contrary, consecutive utterances address the same question under discussion.

Here is a brief informal illustration of how the two principles work together in the derivation of identity of the eventualities in (15): Topic continuity makes the hearer find the strongest common discourse topic for the two utterances, which is the question *What happened to Alena?* in both cases. Exhaustivity requires that each utterance be interpreted exhaustively with respect to that question. That is, the only thing that happened to Alena was breaking her skis, and the only thing that happened to Alena was losing her main means of transport. But since this is the only thing, it must be the same thing. Thus we infer an identity relation between the two eventualities.

Two things are important to note about this approach. On the one hand, it operates exclusively with independently motivated pragmatic principles.

For instance, exhaustivity can be seen as a consequence of Gricean reasoning with the maxim of Quantity restricted by Quality and Relevance: provide as much information as you can without violating Quality, i.e. without making false or unwarranted statements, and without violating Relevance, i.e. without giving irrelevant information (Grice, 1975). Thus, if the speaker utters *Alena broke her skis* addressing the question *What happened to Alena?* and doesn't mention that Alena's car was stolen, this can be indicative of two things: Either the speaker does not consider the theft of Alena's car

relevant, that is, she interprets the question as coming with an implicit domain restriction that excludes Alena's car from consideration, or she does not believe that Alena's car was stolen: $\neg\text{BEL}_S(\text{`Alena's car was stolen'})$. On the assumption that the speaker has an opinion on the question under discussion, this inference can be strengthened to $\text{BEL}_S(\neg\text{`Alena's car was stolen'})$, i.e. the speaker believes that Alena's car was not stolen. This reasoning applied to all the elements in the domain of the question gives us the exhaustivity inference.

Topic continuity can be seen as a manifestation of a more general coherence principle such as *Maximise Discourse Coherence* (MDC) in Asher and Lascarides (2003) or *NEW in Zeevat (2010), which bids you to establish as many and as strong links as possible between utterances in discourse, including anaphoric links, presupposition resolutions, coherence relations, etc.

Notice that from the point of view of cognitive processing, both principles can be viewed as implementations of different aspects of economy. The Principle of Exhaustivity tells the hearer not to add objects to the semantic representation of the utterance beyond those explicitly mentioned. Topic continuity tells one not to add new discourse topics to the structural representation of the discourse, unless the speaker explicitly indicates that a topic change is taking place. Both principles lead towards the construction of the most economical representation in terms of the number of semantic and pragmatic objects involved.

The present approach has the advantage over the relational approach of SDRT, since it operates with fewer basic concepts and principles. The maxims (Gricean or similar) that stand behind exhaustivity are few, self-explanatory, and intuitively sensible. Questions as discourse topics are a relatively well-understood and independently motivated concept, and a coherence maximisation principle that operates just on those is simpler than one that takes all kinds of links, including coherence relations, into account. In contrast, coherence relations are many and no one really knows how many because there are always different possible criteria and degrees of granularity in conceptualising their inventory. In other words, if we managed to provide a pragmatic account of the full variety of semantic relationships between sentences in discourse along the same lines, we would not need to commit to any particular inventory of coherence relations.

On the other hand, since *Elaboration* (or rather identity relations associated with *Elaboration*) is inferred from general pragmatic default principles and need not be triggered by anything specific in the linguistic input, this is a default coherence relation. The question that arises then is how not to infer *Elaboration*. What allows us to sometimes infer other coherence relations? Six papers in this collection, cf. sections 1.3–1.5, deal with various parts of this problem.

Chapter 2 argues that coherence relations like *Narration*, *List*, and *Contrast* are non-default in the sense that they always require some kind of linguistic marking that switches off either the Exhaustivity principle, or the Topic Continuity principle. Chapters 3 and 4 (Jasinskaja and Zeevat, 2008, 2009), introduced in section 1.3, present a

systematic QUD-based analysis of the marking patterns of such relations.

Other coherence relations, most notably *Explanation*, share with *Elaboration* the ability to be inferred without any marking. In order to account for the inference of such relations, the basic pragmatic machinery needs to be extended. This challenge is taken up in chapter 8 (Jasinskaja and Karagjosova, submitted), summarized in section 1.5.

Finally, it is important to note that these two classes of coherence relations—*Narration*, *List*, and *Contrast*, on the one hand, and *Elaboration* and *Explanation*, on the other—correspond roughly to what has become known as coordinating vs. subordinating coherence relations, respectively (Asher and Vieu, 2005). This distinction also has another typical manifestation: Coordinating relations establish flat, non-hierarchical discourse structures. One of the most important factors that influences the prominence of propositions connected in such a non-hierarchical way is *recency*. The most recent proposition introduced in the last sentence is usually more prominent than the earlier ones. It is therefore open for attachment of new discourse material via a coherence relation (Asher and Lascarides, 2003). For example, (18-c) will much rather be understood as an *Elaboration* of (18-b) than of (18-a), i.e. as 'Max did a good job painting the wall' rather than as 'Peter did a good job repairing the roof', because (18-a) and (18-b) stand in a coordinating *Parallel* relation, as they present two similar events.

(18) a. Peter repaired the roof.
b. Max painted the wall.
c. He did a good job.

In contrast, subordinating coherence relations give rise to hierarchical discourse structures and a prominence ranking between main and subordinate material, which to some extent can work against recency. For instance, (19) exemplifies a subordinating coherence relation *Explanation*: (19-b) explains why Peter repaired the roof. The *Elaboration* in (19-c) can attach both to (19-b) and to (19-a), and accordingly allows for both interpretations 'Max did a good job convincing Peter' and 'Peter did a good job repairing the roof'.

(19) a. Peter repaired the roof.
b. Max convinced him.
c. He did a good job.

The generalisation illustrated above is known in the literature as the Right Frontier Constraint (Polanyi, 1988; Asher and Lascarides, 2003), which says that only the last processed discourse node and the nodes it is subordinated to are accessible for attachment of new discourse material.

Even though the papers in this collection concentrate much more on the marking patterns of coordinating vs. subordinating relations, while the prominence effect of this distinction is only touched upon in chapters 8 and 10, it is useful to keep in mind that we are talking about important natural classes of coherence relations. The present

findings about marking patterns contribute to our understanding of coordination vs. subordination in discourse and the proposed theories lay the foundation for future, more comprehensive accounts.

1.3 Marking coordinating relations

While chapter 2 explains how *Elaboration* is inferred by default, i.e. how the identity relation between the eventualities of Alena breaking her skis and losing her main means of transport is inferred in (20-a) by applying the principles of Exhaustivity and Topic Continuity, Zeevat and Jasinskaja (2007, not included in this collection) and the chapters 3 and 4 in this volume (Jasinskaja and Zeevat, 2008, 2009) deal with the question how connectives like *and* and *but* override the *Elaboration* default, that is, for instance, how *and* in (20-b) ensures that Alena breaking her skis and her losing her main means of transport are understood as distinct eventualities, i.e. that skis did not use to be her main means of transport.

(20) a. Alena lost her main means of transport.
She broke her skis.
b. Alena lost her main means of transport.
And she broke her skis.

The solution developed in these papers is based on the idea that connectives like *and* encode *additivity*, that is, that the conjuncts of the connective give *distinct* answers to the same question. This is the difference to the default case: In (20-a), the two sentences give exhaustive answers to the same question (*What happened?*), and the non-distinctness, moreover the identity of the eventualities follows from the combination of exhaustivity and the assumption that the QUD remains the same, as explained in the previous section. Also in (20-b) the QUD remains the same, but here *and* imposes the requirement that the conjuncts give distinct answers, i.e. name distinct eventualities that happened. This is only possible if each answer is non-exhaustive. Thus the distinctness requirement indirectly contradicts the Principle of Exhaustivity. Since constraints encoded by explicit linguistic devices are ranked higher than pragmatic principles, the semantic contribution of the connective overrides exhaustivity.[6]

Chapter 3 develops further the idea that the semantic differences between coordinative connectives are best described in terms of the type of QUD their conjuncts answer. Questions are represented as sets of Hamblin-style alternatives (Hamblin, 1973), e.g. the question *Who smokes?* corresponds to the set of mutually compatible possible answers { *John smokes, Mary smokes, Bill smokes, ...* }. The question types relevant for the description of the connectives differ according to two main parameters: the number and the type of question variables. In terms of the number of variables, the

[6] In optimality-theoretic pragmatics (Zeevat, 2009, 2010) this is modelled by ranking the FAITH constraint over *NEW (≈ Topic Continuity) and RELEVANCE (responsible for exhaustivity effects).

most important distinction is between single and multiple variable questions, which corresponds to the number of dimensions in which the question alternatives differ. The canonical cases are single (*Who snores?*) vs. multiple *wh*-questions, e.g. *Who likes what?*, *Who gave what to whom?*, etc., respectively. The most important variable types are, informally, *wh* for various types of entities that can answer questions like *who*, *what*, *when*, etc., and the *y/n* type for negative vs. positive polarity instantiated by negation and an identity operator of the same logical type. This is the variable type of *yes/no*-questions like *Does John like football?* and corresponds to the word *whether* in embedded questions.

In chapter 3 it is proposed that English *and*, German *und* and the Russian connective *a* just encode additivity and do not impose any constraints on the QUD addressed by the conjuncts. English *but* and German *aber* impose the requirement that their conjuncts give doubly distinct answers to a *wh-y/n*-question—a question with two (or more) variables one of which is a *y/n*-variable and the other is *wh*. That means that the conjuncts must differ along two dimensions: in their polarity and in the instantiation of the *wh*-variable, e.g. *John* vs. *Bill* in (21). The negation need not always be explicit: For instance if the QUD is *Are John and Bill tall?*, *Bill is small* amounts to saying that he is not tall, and in that sense the answer provides a distinct instantiation of the *y/n* variable.

(21) John likes football, *but* Bill doesn't.

(22) John is tall, *but* Bill is small.

A special case of doubly distinct answers to a *wh-y/n*-question are corrections. German has a special connective *sondern* (23) to mark this type of use. Also here, the conjuncts differ in polarity and in the instantiation of a *wh*-variable, e.g. *where*: *in Berlin* vs. *in Paris* in (23). However, specialized correction markers like *sondern* come with an additional presupposition that ultimately leads to the effect that the "wrong" element *Berlin* is "replaced" by the "correct" element *Paris*. The specifics of correction are studied in more detail in chapters 5–7 (Jasinskaja, 2010a, 2012, 2013) and summarized in section 1.4. English does not have a special connective for corrections, so it uses the general *wh-y/n*-marker *but* in this function.

(23) Peter ist nicht in Berlin, *sondern* in Paris.
Peter is not in Berlin but in Paris

Another special case of *wh-y/n* is a *why-y/n*-question, e.g. *Why should and why shouldn't we buy this ring?* in (24). One conjunct of *but* gives a reason for buying the ring, and the other conjunct gives a different reason for not buying it. Again, English uses the general *wh-y/n*-marker *but* in this case, but Russian has a special *why-y/n*-connective *no*. So Russian *no* is used in contexts like (24), but not in contexts like (21)–(23), where the general additive marker *a* is used instead.

(24) This ring is beautiful, *but* expensive.

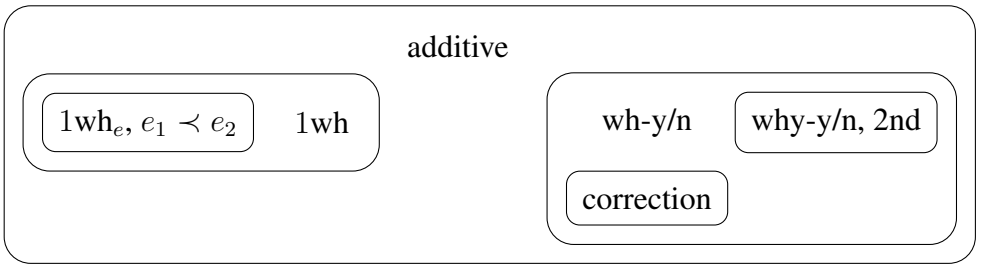

Figure 1.3: Semantic functions of additive and adversative markers

Finally, Russian connective *i* is only used if the QUD has only one *wh*-variable, i.e. the conjuncts only differ along one dimension. Therefore *i* is appropriate in contexts like (25) (What is the weather like?), but not is contexts like (26) (What kind of weather is where?), where the conjuncts differ along two dimensions: the location and the weather state description. In the latter case, the general additive connective *a* is used, whereas in English and in German both functions are covered by the general additive connective *and* and *und*, respectively.

(25) It is snowing *and* the wind is blowing.

(26) In Moscow, it is snowing *and* in Amsterdam, the wind is blowing.

The different functions of the connectives and the relations of specificity between them are summarized in figure 1.3. The figure also contains the function 1*wh*$_e$, $e_1 \prec e_2$ which does not have a dedicated connective in the systems of any of the three languages discussed in chapter 3, but is interesting from the point of view of coherence relations. The function is a special case of a single variable *wh*-question (1wh), but restricts it to questions about events (1wh$_e$), e.g. *What happened?* In addition, there is a requirement that the events answering the question are not only distinct, but also stand in a temporal precedence relation ($e_1 \prec e_2$). These constraints constitute the conditions for the coherence relation of *Narration*, as in (27). Russian uses the 1*wh*-marker *i* in this case, whereas English and German use the general additive markers *and* and *und*.

(27) She jumped on the horse *and* rode into the sunset.

Chapter 4 (Jasinskaja and Zeevat, 2008) discusses a number of more tricky cases, including topic change and parenthetical uses of the Russian *a*, as well as causal uses of the Russian *i*, much discussed in Russian linguistics, which, however, are all shown to fall under the basic schemata discussed above.

A further ingredient of Jasinskaja and Zeevat's analysis is the mechanism of *blocking*. It assumes that the marking of the features in figure 1.3 is more or less obligatory when an appropriate marker is available. In contexts where both a more general and a more specific marker is applicable, the more specific one is preferred. Therefore the

usage of more general markers is not as broad as their weak semantics suggests, but is restricted by the presence of more specific markers in the system. For instance, the German *aber* encodes *wh-y/n* and therefore is compatible with correction, which is a special case of *wh-y/n*. However, it is blocked by *sondern* and is therefore not normally used in contexts like (23). Similarly, Russian *a* is blocked "on both sides" by *no* in *why-y/n* and by *i* in 1*wh*. As a result, it is restricted to opposition uses like (21) and (22), correction (23), and pair list answers to double (or multiple) *wh*-questions like (26). If we break up the hierarchically organized functions in figure 1.3 into a set of complementarily distributed uses that result from blocking, the semantic space of the coherence relations based on additivity can be represented as in figures 1.4–1.6, which also show how that space is carved up by the English, the German and the Russian connective systems, respectively.[7]

Thinking in terms of semantic spaces reveals the shortcomings of the traditional definitions of coherence relations. For instance, in SDRT (Asher and Lascarides, 2003) *Narration* is defined as a relation connecting descriptions of events occurring in a sequence (1wh_e, $e_1 \prec e_2$), such as those in (27). *Contrast* is defined to include opposition like (21) and (22) and argumentative contrast (24), and corresponds to our *wh-y/n*. The remaining functions (1*wh*′ and *wh-wh*) fall roughly under the definition of *Parallel*. As can be seen in figure 1.7, the resulting system neither reflects the system of any particular language among the three discussed in chapter 3, nor is it cross-linguistically valid. In fact, a closer look at existing definitions of *Contrast* (Kehler, 2002; Asher and Lascarides, 2003) suggests that they were rather influenced by the anglocentric perspective and the semantics of the connective *but*. The ideas behind the definition of *Narration*, however, have little to do with the semantics of connectives, but rather with common pragmatic inferences based on principles such as the Gricean *be orderly* (Grice, 1975) and iconicity in the chronological presentation of event sequences (Jakobson, 1971). In other words, the existing inventories of coherence relations are based on rather heterogeneous considerations and the ultimate result, such as that in figure 1.7, appears rather arbitrary.

A more principled approach would be either to follow systematically the categories encoded in a particular language (as was previously advocated by e.g. Hovy and Maier, 1995, cf. related discussion in chapter 8), and thereby commit to language-specific sets of coherence relations, or to adopt a maximally fine-grained classification, where a different coherence relation corresponds to every single node in the semantic space. This would result in a cross-linguistically valid inventory of relations.[8] However, on this

[7]The figures represent semantic maps that are introduced in more detail in chapter 5 (Jasinskaja, 2010a) and chapter 8 (Jasinskaja and Karagjosova, submitted) in this collection. The function *wh-y/n*′ corresponds to *wh-y/n* blocked by *why-y/n*; *wh-y/n*′′ corresponds to *wh-y/n* blocked by both *why-y/n* and *correction*; 1*wh*′ is 1*wh* blocked by 1wh_e, $e_1 \prec e_2$; finally, *wh-wh* (multiple *wh* questions) is what is left over of a general additive marker if both 1*wh* and *wh-y/n* questions are excluded.

[8]As discussed in more detail in chapter 5 (Jasinskaja, 2010a), the semantic space is figures 1.4–1.7 is a refinement of the semantic spaces developed by Malchukov (2004) and Mauri (2008), which are based on much more representative language samples than just the three languages discussed in chapter 3.

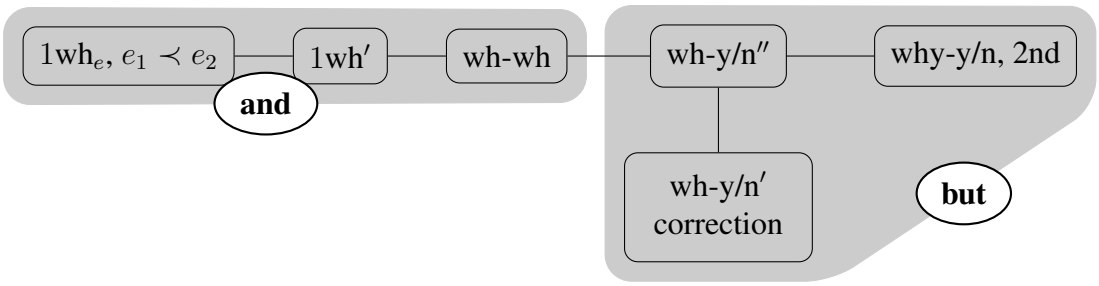

Figure 1.4: English *and* and *but*

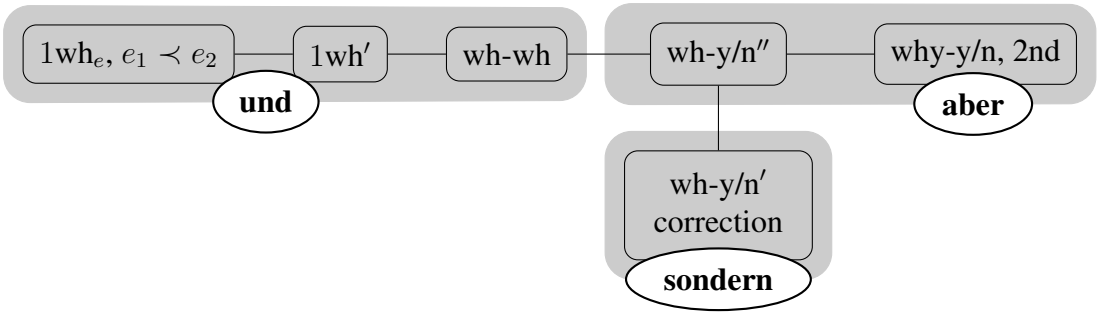

Figure 1.5: German *und*, *aber*, *sondern*

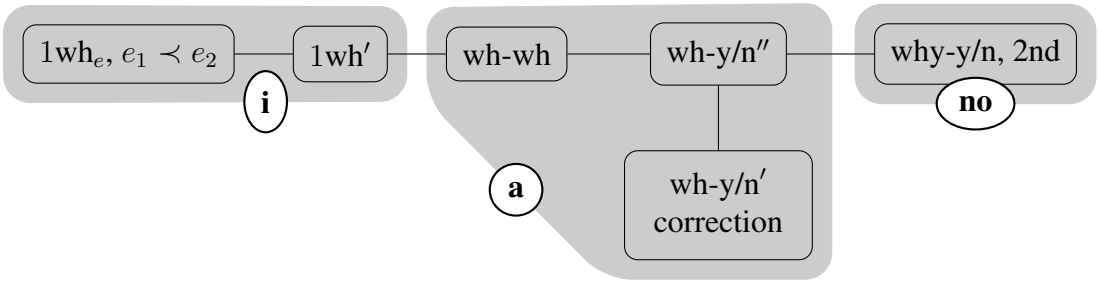

Figure 1.6: Russian *i*, *a* and *no*

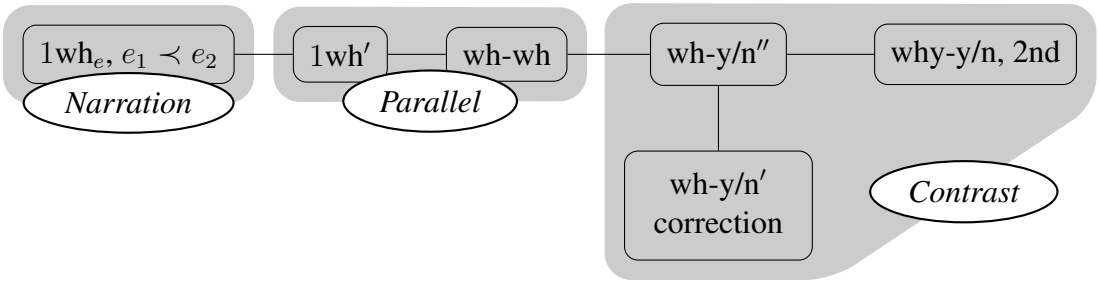

Figure 1.7: The semantic space of coherence relations *Narration*, *Parallel*, *Contrast*

approach one could just as well not define any coherence relations at all, but operate directly with the proposed QUD categories 1*wh*, *wh-y/n*, etc. In other words, QUDs turn out useful not only as constraints on default pragmatic inferences, but also as a basis for a cross-linguistic semantics of connectives that encode non-default relations.

1.4 Correction

The next three chapters 5–7 (Jasinskaja, 2010a, 2012, 2013) are concerned with a single construction, which expresses *correction*. In German, correction is unambiguously signalled by the connective *sondern*, already mentioned in the previous section and illustrated in example (23), repeated below:

(28) Peter ist nicht in Berlin, *sondern* in Paris.
Peter is not in Berlin but in Paris

This construction is particularly interesting because it combines properties of coherence relations of both kinds discussed in the previous two sections. On the one hand, it is similar to *Contrast* relations in that it is marked in many languages by adversative markers like *but* as in (29), cf. section 1.3. On the other hand, it resembles *Elaboration* in that it can be inferred without any special marking as in (30), and it has the characteristic *Elaboration* semantics since the two sentences in (30) or the two (elliptic) clauses in (29) refer to the same situation, cf. section 1.2.

(29) Peter is not in Berlin, *but* in Paris.

(30) Peter is not in Berlin. He is in Paris.

Correction in the present sense has not attracted much attention within the relational approach to discourse, but it does present a puzzle for that approach. Traditionally, coordinating and subordinating coherence relations are viewed as distinct classes. A relation cannot be coordinating and subordinating at the same time (see e.g. Txurruka, 2003). However, *Contrast* is usually considered a coordinating relation, and *Elaboration* a subordinating one. It seems that to work around this conflict the relational approach would require treating each occurrence of correction either as an instance of *Contrast*, thereby ignoring its elaborative properties, or as an instance of *Elaboration*, thereby ignoring its contrastive properties. In fact, Asher (2013) discusses examples similar to (30), assuming without discussion that the coherence relation there is simply *Elaboration*. But then, how do we explain the strong intuition that (29) expresses the same relation, although it is marked by *but*, which signals *Contrast*?

Chapters 5–7 develop a QUD-based account of various properties of corrections. In contrast to the relational approach, there is no pressure to squeeze those properties into predefined bundles, and the various aspects of correction can be considered one at a time. Chapters 5 and 6 are concerned with the contrastive side of correction, whereas chapter 7 is concerned with its elaborative side.

The chapters further develop the approach of chapter 3 (Jasinskaja and Zeevat, 2009), where the conjuncts of connectives like the German *sondern* are a special case of doubly distinct answers to a *wh-y/n*-question. The distinct instantiations of the *y/n*-variable manifest themselves in the fact that one conjunct of *sondern* is always negative, while the other is positive. The distinct instantiations of the *wh*-variable are the corrected and the correcting part of the utterance, e.g. *Berlin* vs. *Paris* in (28) and (29). Since *but* is a marker of *wh-y/n*, this approach explains why *but* can be used to mark correction. However, it remains less clear what constitutes the specifics of correction and how it is different from other kinds of contrast that involve explicit negation, such as (31). After all, in both cases the conjuncts have opposite polarity and distinct participants.

(31) John doesn't like football, *but* Bill does.

Chapter 5 concentrates on the information-structural differences between corrections like (29) and contrastive pairs like (31). While in both cases a *wh-y/n*-question is addressed, (29) and (31) differ in the way they split the double variable question into single variable subquestions. In (31), the speaker 'goes by people', i.e. splits the *wh-y/n*-question into a series of *yes/no*-questions about each relevant person in the domain: *Does John like football?*, *Does Bill like football?*, etc. This leads to the information-structural pattern where *John*, *Bill*, etc. is the contrastive topic, whereas positive vs. negative polarity of the clauses is the focus (Büring, 2003, see also Lohnstein 2016 for an overview on polarity and verum focus). In (29), the speaker 'goes by polarity', i.e. splits the *wh-y/n*-question into constituent questions *Where is Peter not?* and *Where is Peter?* This leads to the prediction that the constituents answering the *where*-question, *in Berlin* and *in Paris*, constitute the foci in (29) and bear the nuclear accent. Extending Büring's (2003) theory of contrastive topic to *wh-y/n*-questions, this would also mean that the polarity values constitute the contrastive topics. However, we do not see any contrastive topic accents, let alone on polarity, in corrections like (29). Chapter 5 argues that this is due to a general tendency of non-individuated semantic objects, such as polarity values, to avoid the topic role.

The second way in which corrections differ from contrastive pairs is what Jasinskaja and Zeevat (2009) call the 'correction presupposition'. The chapters develop the view that this 'presupposition' consists in the 'replacive' property of corrections, in the sense of Jacobs' (1982) replacive negation. That is, one clause of a correction negates a proposition, whereas the other clause 'replaces' the 'wrong' part of that proposition by a 'correct' element. Chapter 7 argues that this replacive property is a consequence of the *Elaboration* semantics of correction—the fact that the two clauses describe the same situation, which in turn is derived along the general schema for *Elaboration* relations, proposed in chapter 2, i.e. as a consequence of applying the principles of *Exaustivity* and *Topic Continuity*. The assumption is that both clauses or sentences in a correction address the same question, such as *Where is Peter?* in (29) and (30). The negative clause negates an exhaustive answer to this question: it is not in Berlin that

Peter is. The positive clause gives the correct exhaustive answer *Paris*. This leads to the effect that both clauses are talking about the same state of Peter being in a place, and the wrong location is replaced by the correct one for that state.

Finally, it turns out that there exist non-trivial semantic differences among corrections, which manifest themselves in particular in their *symmetric* vs. *asymmetric* behaviour, studied in detail in chapter 6. For instance, the German *sondern* and the English *but* are asymmetric. *Sondern* is fine if the negative conjunct precedes the positive one, but infelicitous with the opposite order, cf. (32). The English *but* is possible with the positive-negative order, but does not encode correction in the sense of replacivity: (33) does not express the idea that John is in Paris rather than in Berlin, but the idea that John is only in Berlin, while he could have been in both places, which is rather unlikely in the literal sense of being in a place.

(32) *Peter ist in Paris, *sondern nicht* in Berlin.
Peter is in Paris CORR not in Berlin

(33) ?John is in Berlin, *but not* in Paris.

Other correction markers show symmetric behaviour. For instance, the Russian *a* expresses correction regardless of the order of conjuncts:

(34) a. Oleg *ne* v Pariže, *a* v Berline.
Oleg not in Paris CORR in Berlin
Oleg is not in Paris, but in Berlin.

b. Oleg v Berline, *a* *ne* v Pariže.
Oleg in Berlin CORR not in Paris
Oleg is in Berlin, and not in Paris.

Similarly, unmarked correction in (30) is symmetric. The replacive relationship remains intact even if the order of sentences is reversed.

The explanation proposed in chapter 6 relates the asymmetry of *but* to the feature 2ND, cf. figures 1.4–1.6, encoded by this connective. The feature consists in the requirement that the 2nd conjunct must provide a more resolving answer to the current *question under dispute*, than the 1st conjunct. The question under dispute is a specific variety of the question under discussion, with respect to which the communication participants hold distinct positions, i.e. their belief states entail distinct answers to that question. In the argumentative uses of *but*, such as (24), repeated below, the question under dispute is whether we should buy this ring. The first conjunct is an argument for buying the ring, while the second is an argument against it. Since *but* conventionally encodes that the second conjunct is more resolving, the whole of (35) is understood rather as a suggestion not to buy the ring. In contrast, if the order of conjuncts is reversed, as in (36), the argument for buying the ring is stronger.

(35) This ring is beautiful, *but* expensive.

(36) This ring is expensive, *but* beautiful.

In corrections like (29), the question under dispute is where Peter is. An exhaustive positive answer to this question is always more resolving than the negative one. The answer 'Peter is in Paris' picks out a single cell of the question partition, suggesting that Peter is in Paris and nowhere else. The negative answer 'Peter is not in Berlin' only excludes Berlin from the possible places Peter could be in, but leaves it open where he actually is. Therefore the positive answer must come second.

In contrast, the Russian *a* does not encode the feature 2ND, which manifests itself not only in its symmetric behaviour in corrections, but also in its non-resolving character in argumentative contrast contexts like (35) and (36) (cf. especially chapter 4). In unmarked corrections, in turn, there is nothing that could encode anything at all, so again, no asymmetry is encoded.

In sum, these studies show that the QUD approach can be used not only to explain why corrections can be marked by the *Contrast* marker *but*, but also how they differ from other uses of *but*: why they have a different information structure and a different semantics. Moreover, the notion of QUD can be used to encode the differences between symmetric and asymmetric correction markers and to explain why unmarked corrections are symmetric. It is hard to see how we could have arrived at this sophisticated picture using the vocabulary of coherence relations such as *Contrast*, *Elaboration*, and the like. The problem with coherence relations is that they represent 'packages' of features, which perhaps often, but not always co-occur. For instance, the '*Contrast* package' includes coordination and, in SDRT terms, the requirement that the units it connects be maximally opposite'. The '*Elaboration* package' includes subordination and mereological part-whole and identity relations between the described situations (Asher and Lascarides, 2003). But for corrections, it seems that we do not want to choose one of the two, or to buy a whole package. Instead, we can use QUDs, exhaustivity, topic continuity, the notion of splitting a question into subquestions and the notion of question resolution to characterize the contribution of each particular ingredient: the exhaustivity and topic continuity for the *Elaboration* semantics of corrections and for unmarked corrections, *wh-y/n*-questions for the contribution of correction markers, and questions under dispute for the specific asymmetric property of *but*. Under this approach, it is not surprising that in terms of marking patterns, corrections sometimes behave like subordinating and sometimes like coordinating relations. This work, however, does not address at all the question of whether corrections should be considered coordinating, subordinating, or both in terms of their effect on prominence (see discussion in section 1.2), which remains a question for future research.

1.5 Grounding and acceptance: Explaining subordinating relations

The approach to the inference of coherence relations proposed in chapter 2 (Jasinskaja, 2010b) has one important problem. Despite all the restrictions on the default inferences that come from explicit connective devices studied in subsequent papers, the theory in chapter 2 still overgenerates *Elaborations*. Let us first briefly summarize what we have reached so far.

As explained in chapters 2 and 7, the general pragmatic principles of *Topic Continuity* and *Exhaustivity* lead to a default inference of *Elaboration* (37), and *Correction* as a special case of *Elaboration* (38). That is, the two sentences in (37) and (38) describe the same event of John praising someone. The same inference takes place in (39): *Topic Continuity* makes us assume that the two sentences address the same question *Who did John praise?*; *Exhaustivity* makes us interpret each sentence exhaustively with respect to that question. That is, the only person praised by John is Bill, and the only person praised by John is Mary. But this is only possible if Bill and Mary are the same person, which is quite unlikely under normal circumstances. This is the reason why (39) sounds odd when uttered with the same 'full stop intonation' on each sentence as (37) and (38).

(37) John praised a student. He praised Mary.

(38) John didn't praise Bill. He praised Mary.

(39) ??John praised Bill. He praised Mary.

It is easy to repair (39) by inserting connective devices such as *also* (40) and *and* (41), or simply by replacing fullstops by commas, or 'full stop intonation' by 'continuation intonation' (42). All these devices signal in one way or another that either *Topic Continuity* or *Exhaustivity* is not applied (see Zeevat and Jasinskaja, 2007, and chapters 3 and 4 in this volume).

(40) John praised Bill. He *also* praised Mary.

(41) John praised Bill *and* he praised Mary.

(42) John praised Bill, he praised Mary...

However, there is a range of other coherence relations, most notably *Explanation* (43) and *Background* (44),[9] that can be unmarked, i.e. expressed without any connective devices that switch off *Exhaustivity* or *Topic Continuity*, and with the same intonation as the examples (37)–(39):

[9]In an *Explanation*, the second sentence 'explains' the first, by giving causes, reasons, evidence, motivation, etc. In a *Background* relation, the second sentence provides a 'backdrop' for the first one. In particular it can introduce objects talked about in the first sentence, e.g. *Bill* in (44). See chapter 8 (Jasinskaja and Karagjosova, submitted) for an overview of coherence relations.

(43) John praised Bill. He wrote the best term paper.

(44) John praised Bill. Bill is our new PhD student.

The problem is that the theory in chapter 2 predicts *Elaboration* in these cases, too. For instance in (43), the pragmatic principles make us interpret each sentence exhaustively with respect to their shared question *What happened?* That implies that the only relevant event is that John praised Bill. And the only relevant event is that Bill wrote the best term paper. Again, this would only be possible if John praising Bill and Bill writing the paper were the same event. But this is obviously not the case since these events occur at distinct times, most probably at distinct locations, and involving distinct sets of participants. In other words, we derive the same kind of contradiction as we did in (39).

Chapter 8 (Jasinskaja and Karagjosova, submitted)[10]—addresses this problem by substantially extending the underlying pragmatic framework. Putting it in the terminology we have used so far, the proposal implies relaxing the *Topic Continuity* principle. In order to satisfy *Topic Continuity* it is no longer necessary that adjacent utterances U_1 and U_2 address literally the same QUD. It is enough if the communicative goal of U_1 is not discarded while U_2 is processed, that is, the communication participants still address directly or indirectly the QUD of U_1 while they deal with U_2. Since communicative goals are discarded once they are reached, i.e. QUDs are discarded once they are fully resolved, U_1 would be normally kept on hold during the processing of U_2 only if U_1's goal cannot be reached before that of U_2 is reached.

Similar notions of relationships between communicative goals have been proposed before (e.g. *dominance* in Grosz and Sidner, 1986). The main contribution of chapter 8 is in the specific view of what can prevent reaching a goal, and in what sense subordinating relations like *Elaboration*, *Explanation* and *Background* serve to solve the problem that caused the interruption in the processing of the preceding utterance. The central idea is that before the content of an utterance enters common ground, it must be *grounded*—heard, as well as acoustically and semantically understood—by the hearer and it must be *accepted*—e.g. in order for a proposition to become part of the shared beliefs, the hearer must agree to believe it (Clark, 1996; Ginzburg, 2012). If an utterance cannot grounded or accepted, its goal is not yet reached and an *Elaboration*, an *Explanation*, or *Background* must be provided, to make the grounding or acceptance of the utterance possible. For example, in (44) above the speaker anticipates that the hearer might not be able to construct a well-formed semantic representation of the first sentence because he might not know who the proper name *Bill* refers to. This is a grounding problem at the level of semantic understanding and is solved in (44) by pro-

[10]Jasinskaja and Karagjosova (submitted) is intended as a chapter for Wiley's *Companion to Semantics*, and therefore sets out the landscape for coherence relations, including their definitions, linguistic effects, and criteria for choosing a particular inventory of coherence relations in a broader perspective. It has therefore been cited at various places above for reference to these general concepts. The paper, however, also contains a case study on the notion of discourse-structural subordination, which is summarized in this section.

viding *Background* information. Problems of reference resolution, as well as problems of lexical access can also be helped by providing a *Reformulation*, a special case of *Elaboration*. And problems of acceptance can be helped e.g. by providing *Evidence*, a special case of *Explanation*.

Crucially, in the strong version of this theory, failure to get U_1 grounded or accepted is the speaker's all and only reason to interrupt its processing and to keep its communicative goal or QUD on hold. As long as the communicative goal of U_1 is on hold, the principle of *Topic Continuity* is not violated and there is no need to use any linguistic devices to signal that the speaker is moving on to a different QUD. For the hearer that means that when he hears a sequence of utterances $\langle U_1, U_2 \rangle$ and there is no connective like *and*, no particle like *also*, no continuation intonation, cf. (40)–(42), no contrastive topic marking, etc. to indicate that the speaker is moving on to the next point on the agenda, then the hearer can assume that the QUD of U_1 is still open, the speaker must be dealing with some potential grounding or acceptance issue, and therefore the relation between the utterances must be one of the broad class of *Elaboration*, *Explanation* or *Background*. In other words, this theory is not as strong as that in chapter 2, which used to see *Elaborations* all over the place, but is it still rather strong because it implies that all subordinating coherence relations that admit this unmarked pattern of expression should lend themselves to an analysis in terms of remedy for a grounding or an acceptance failure. To show that this is possible remains a challenge for future research.

Moreover, chapter 8 shows a way to account for the influence of subordinating relations on prominence, mentioned briefly in section 1.2. The state of a communicative goal or a QUD of being "on hold" has been captured in existing goal-based and QUD-based approaches to discourse using the mechanism of a goal/QUD *stack* (see Grosz and Sidner, 1986; Roberts, 1996; Ginzburg, 2012). Superordinate goals or QUDs are stored for future reference on the stack underneath the subordinate ones. The topmost goal/QUD on the stack is the one the conversation is dealing with at the current moment. Once the topmost QUD is resolved it is popped off the stack and the superordinate QUD lying underneath it becomes topmost again. Assuming that semantic objects associated with the topmost goal or QUD are more prominent than the others, and that semantic objects associated with a goal/QUD on the stack lose their prominence less rapidly than those that are not on the stack any more, this explains why superordinate material remains accessible for future discourse attachment and reference, as explained in section 1.2. In other words, stacks in a goal-based model capture the same generalization as the Right Frontier Constraint does on the basis of discourse trees in a relational model.

Finally, notice that the theory developed in chapter 8 is based entirely on independently motivated categories and principles. Grounding and acceptance are central concepts of dialogue semantics which manifest themselves ubiquitously in the shape of backchannelling signals, clarification requests, and other kinds of negotiation of the common ground (Traum, 1994; Clark, 1996; Rodríguez and Schlangen, 2004;

Ginzburg et al., 2007; Ginzburg, 2012). Reusing these concepts to explain the notion of discourse-structural subordination is a step forward with respect to the current state in the relational approach. In earlier relational theories subordinating and coordinating relations were simply listed. Why, for instance, *Explanation* should be subordinating and *Contrast* coordinating remained unexplained. In a more recent version of SDRT, Asher and Vieu (2005) provide a more flexible view of the subordination/coordination distinction. However, they define this distinction ultimately in terms of its linguistic manifestations, and give up on the idea of giving it a general semantic definition. In contrast, we define discourse-structural subordination in terms of independently motivated concepts of dialogue semantics in a goal-based framework, in a way that is able to explain both the marking patterns of subordinating vs. coordinating relations and their effect on prominence.

Let me briefly summarize what has been done in the last four sections. The seven papers presented in sections 1.2–1.5 develop an account of coherence relations in a general pragmatic framework. The relations do not have any special status in the theory, but are epiphenomena of reasoning with Gricean maxims, the principle of *Topic Continuity*, and the practices for negotiating common ground in dialogue. In this approach, coherence relations boil down to inferences of the same sort as, for instance, Gricean implicatures, and to relations encoded by linguistic markers like *and*, *but*, etc. In contrast to more widely familiar kinds of Gricean implicature, coherence relations are 'relational implicatures', representing meaningful links between two or more speech acts. The next section takes a look at more canonical kinds of pragmatic inference—conversational and conventional implicature within the QUD-based approach.

1.6 Further applications of the QUD approach

One of the goals of the theories of discourse coherence is to describe and to model the 'extra content' that distinguishes a coherent discourse from an arbitrary sequence of sentences. In relational theories that means primarily to have semantic definitions of coherence relations like *Elaboration*, *Explanation*, *Contrast*, *Narration*, etc., and a theory of how that content is encoded or implicated. In a goal-based or QUD-based approach, the same task boils down to developing a theory of how this extra content results from the goal/QUD structure of discourse and some general pragmatic or processing principles. The papers summarized in the previous sections address various parts of this task.

Another goal of the theories of discourse coherence is to explain context-sensitive phenomena in language. Motivation for a specific theoretical approach is often sought in how well it is able to account for the relevant linguistic phenomena. Within the relational approach, a lot of effort has been directed at showing that the distinctions between the postulated types of coherence relations are relevant for an account of e.g. tense, ellipsis, pronoun resolution, lexical disambiguation, presupposition, implica-

ture, etc. (Lascarides and Asher, 1993; Kehler, 2002; Asher and Lascarides, 1995, 1998b; Asher, 2013). In a QUD-based approach, the goal is to show that QUDs provide a useful theoretical tool to account for these phenomena.

The last two chapters in this volume pursue this goal and present applications of the QUD-based notion of discourse structure to resolve widely discussed problems related to conversational and conventional implicatures. Chapter 9 (Jasinskaja et al., 2017) studies the phenomenon of discourse-level exhaustivity implicature, arguing that sequences of multiple sentences can give rise to exhaustivity implicatures that cannot be reduced to a sum of the implicatures of individual sentences and that a QUD-based notion of discourse structure is needed in order to be able to account for such inferences. Chapter 10 (Jasinskaja, in revision) takes up the issue of variable (not-)at-issue status of appositive relative clauses and explains it in terms of the distinction between subordinating and coordinating coherence relations, its influence on prominence and accessibility for discourse attachment and anaphora, whereas QUDs play a central role in establishing a link between discourse structure and the QUD-based notion of at-issueness proposed by Simons et al. (2011).

Discourse-level exhaustivity

The issue of discourse-level implicatures came up as a side issue in the recent fierce debate, known as the localist-globalist debate, on the question whether implicatures, such as quantity implicatures, including exhaustivity, are generated *globally*, i.e. by Gricean reasoning (illustrated on p. 17 for example (15) in section 1.2) no lower than at the level of a whole speech act, or *locally*, i.e. by a special exhaustivity operator that may appear at various levels in the sentence structure. The main argument of the localists is the existence of so called "embedded implicatures", i.e. cases which can arguably only be accounted for if we assume that the pragmatic strengthening that we thought was brought about by Gricean reasoning occurs in the scope of a semantic operator in the sentence structure. An example similar to those brought up as evidence for embedded implicatures is discussed in chapter 7, where I argue that negation in a correction, e.g. in the first sentence of (45), must be interpreted as taking scope over exhaustivity. That is, it is not the case that John praised Bill and no one else, which is consistent with the continuation that John praised Bill and Mary.

(45) John didn't praise BILL. He praised BILL and MARY.

The classical Gricean approach seems to fail on such examples. The global implicature of the first sentence would be, depending on the assumed set of alternatives, either that John praised everyone else except Bill, or no implicature at all, i.e. the exhaustive interpretation is identical with the literal meaning: it is not the case that John praised Bill. Both versions are inconsistent with the continuation in the second sentence.[11]

[11]This is not to suggest that examples like (45) present particularly strong evidence against the Gricean approach, since there are ways to give them an independent explanation. In chapter 7, for

Numerous examples of this and other kinds have led, for instance, Chierchia (2004) and Fox (2007) to drop altogether the idea that upper-bounded construals like *Bill* $\rightsquigarrow$ *Bill and no one else*, *three* $\rightsquigarrow$ *three and not more*, *some* $\rightsquigarrow$ *some but not all* are pragmatic inferences, and to adopt the view that they are defeasible entailments generated by silent exhaustivity operators in the sentence structure during the compositional interpretation of a sentence.

There followed a number of globalist replies (Sauerland, 2004; van Rooij and Schulz, 2004; Russell, 2006) providing alternative explanations of the localist counterexamples, or explaining them away as effects of independently motivated processes. This part of the argument was directed at showing that globalist accounts can handle apparent embedded implicatures, so we can do without exhaustivity operators in the sentence structure. The second part of the argument, presented most systematically by Geurts (2010), was to show that we cannot do without Gricean inferences to derive the upper-bounded construals. Discourse-level quantity implicatures is the case in point. Geurts (2007, 2010) illustrates this point with examples like (46). The one-sentence answer (46-a) and the two-sentence answer (46-b) give rise to the same exhaustivity inference: Cleo visited Naples, Rome, and Ravenna, *and no other places*.

(46) Tony: Which places did you see on your trip to Italy?
Cleo: a. I went to Naples, Rome, and Ravenna.
b. Julius and I first went to Naples and Rome together. Then, while he went to see Milan, I visited Ravenna.

The two-sentence answer is problematic for the localist account. It would require the exhaustivity operator to take scope over two sentences, but there is no syntactic structure above the level of a single sentence, so there is no structure into which such an operator could be introduced. In contrast, the Gricean approach provides a most straightforward solution. The input to Gricean reasoning is a speech act ($\approx$ one sentence). However, it is a common view in discourse semantics that a discourse is itself a complex speech act that consists of simple speech acts (see e.g. Asher and Lascarides, 2003). On this view, one should expect Gricean reasoning to apply to single-sentence speech acts and multi-sentence speech acts alike, so discourse-level implicatures should be generated in much the same way as sentence-level implicatures.

Chapter 9 (Jasinskaja et al., 2017) in this volume puts this argument to a test. First, it argues that many apparent discourse-level exhaustivity implicatures, including the one in (46), can be reduced to the sum of the implicatures of the individual sentences, and therefore do not present a problem to the localist approach. However, Jasinskaja et al. also point out cases in which such a reduction is not possible, so the Gricean account predicts a discourse-level implicature whereas the localist account does not. In two experiments it is shown that indeed upper-bounded inferences are drawn by the participants in the critical cases, which supports the Gricean approach.

instance, the mechanism of metarepresentational negation is used to deal with such cases. There exist more compelling arguments for embedded implicatures, but their discussion would take us too far afield.

Second, Gricean reasoning is reasoning with alternatives. It draws inferences from what the speaker said and what she could have said instead but didn't. What this set of alternatives is like is a challenging research question that has generated a lot of discussion (see e.g. Rooth, 1992; Zeevat, 2004; Katzir, 2007; Geurts, 2010; Fox and Katzir, 2011). However, a widely accepted view is that at least one of the factors that determines the set of alternatives used in Gricean reasoning is the information structure of the sentence. The set is generated by replacing the focused constituent (the constituent bearing the nuclear pitch accent, indicated by small caps in (47)) by values of the same semantic type (cf. Rooth, 1992). Thus, the set of alternatives for (47-a) is 'John praised Bill', 'John praised Sue', 'John praised Peter', etc., and the corresponding exhaustivity implicature is 'John praised Mary and no one else', whereas the set of alternatives for (47-a) is 'Bill praised Mary', 'Sue praised Mary', 'Peter praised Mary', etc., and the corresponding exhaustivity implicature is 'John and no one else praised Mary'.

(47) a. John praised [MARY]$_F$.
b. [JOHN]$_F$ praised Mary.

However, if we want to apply Gricean reasoning above the sentence level then we should also have a notion of alternatives above the sentence level. Once again, the information-structural partition into focus and background does not apply above the sentence level, so it is less clear how to generate alternative sets for sequences of two or more sentences. While different takes are possible on this matter, it seems that the most straightforward solution, proposed in chapter 9, is to use directly the Hamblin alternatives of the QUD dominating the sequence. However, this presupposes a QUD-based discourse structure, which associates questions and respective alternative sets to discourse nodes at all structural levels. In other words, the need to account for discourse-level implicatures provides further motivation for QUD-based discourse structure.

Coming back to the globalist/localist issue, more recent studies (see e.g. Gajewski and Sharvit, 2009; Chemla and Spector, 2011; Sauerland, 2012) have contested many of the arguments presented by the globalist camp. There is growing consensus that exhaustivity operators in the sentence structure are unavoidable after all. However, the argument related to discourse-level implicatures has not been contested. That means that, as far as we can tell, Gricean quantity-based reasoning (or some comparable pragmatic process) is also unavoidable. In sum, the study in chapter 9 presents an example of how an argument for QUD-based discourse structure contributes to a debate on a general pragmatic issue of the status of Gricean quantity implicatures.

Variable (not-)at-issue status

The last chapter 10 in this volume (Jasinskaja, in revision) addresses another hot issue in pragmatics. It is the issue of the pragmatic status of secondary, peripheral, or as they are nowadays called *not-at-issue* components of sentence content.

In his influential book, Chris Potts (2005) revived the notion of *conventional implicature* but refurbished it to cover a different set of phenomena than those originally intended by Grice. Unlike presuppositions, Potts' conventional implicatures are not backgrounded, i.e. normally introduce new information rather than referring to old information in discourse. Unlike conversational implicatures, they are *conventional*, i.e. associated with specific lexical items, morphological markers, or grammatical constructions, which include expressive language (*damn, jerk*), honorifics, parentheticals, nominal appositives, and non-restrictive, or appositive relative clauses. The focus of chapter 10 is in particular on appositive relative clauses, henceforth ARCs.

Two properties of conventional implicatures, which they, by the way, share with presuppositions, are of particular interest: projection and non-challengeability. Projection refers to the ability of certain parts of sentence content to be interpreted outside the scope of semantic operators in whose scope they appear to be syntactically. For instance in (48) both the proposition 'John is smart' and the proposition 'John is competent' appear in the syntactic scope of *doubt* (example borrowed from Schlenker, 2013). However, in (48-a) 'John is smart' is expressed by an ARC and projects: The speaker does not doubt that John is smart. In contrast, in (48-b) the speaker doubts that John is smart and competent.

(48) a. I doubt that John, who is smart, is competent. $\rightsquigarrow$ John is smart
b. I doubt that John is smart and competent. $\not\rightsquigarrow$ John is smart

The second property is the inability of conventionally implicated content to be *challenged*, i.e. directly rejected by utterances like *No, that's not true!* or *No, he/she isn't*. For example, in (49-a) there is a strong preference for the rejection to be interpreted as targeting the main clause. In (49-b), the rejection is altogether infelicitous because it does not fit the main clause and the relative clause is un-challengeable (example adapted from Amaral et al., 2007).

(49) a. A: John, who is smart, is competent.
B: No, he isn't. $\rightsquigarrow$ John isn't competent.
b. A: Edna, who is a fearless leader, started the descent.
B: #No, she isn't.

The view of conventional implicature (and presupposition!) as something conventional has been challenged in a series of papers by David Beaver, Craige Roberts, Mandy Simons and Judith Tonhauser (see e.g. Simons et al., 2011). They see projection and non-challengeability primarily as properties of *not-at-issue content*, which receives in their theory a radically pragmatic definition. Somewhat short-cutting the details, the definition says that a proposition is at-issue iff it is relevant to the current QUD. But since the QUD is given by the context, and the current QUD changes from utterance to utterance in discourse, this definition amounts to saying that projection and non-challengeability are highly context-sensitive phenomena, far from being preset in the lexicon or the grammar for specific linguistic units. For example, Simons et al. (2017)

argue that the projection behaviour of the complements of factive verbs like *know* and *regret* depends on the context represented by the QUD, contrary to the well-established assumption that it is lexically determined.

A major challenge to the conventionalist view of, specifically, the non-challengeability property of ARCs comes from the observation made by AnderBois et al. (2011) and confirmed experimentally by Syrett and Koev (2015) that the ability of ARCs to be directly rejected is dependent on their position in the sentence. Sentence-final ARCs are easier to reject than sentence-medial ones, as shown by the following example from AnderBois et al. (2011):

(50) A: His husband, who had prostate cancer, was being treated at the Dominican Hospital.
B: ??No, he had lung cancer.

(51) A: He took care of his husband, who had prostate cancer.
B: No, he had lung cancer.

On the face of it, it looks like speaker B's rejection simply tries to target the most recent clause. This simple generalization is at the core of Koev's (2013) theoretical analysis of this phenomenon. But if this is recency, and recency, as we know, is a factor that affects prominence of discourse referents (cf. discussion in section 1.2), then we should be wondering whether challengeability and at-issue status of an ARC directly depend on the prominence of its content in discourse. If this is right, then all we know about discourse prominence should, in principle, apply to challengeability and at-issue status as well.

This is the central question studied in chapter 10. The chapter explores the multifaceted consequences of this generalization. In the context of the view of discourse structure developed in my earlier work (summarized in sections 1.2–1.5), the first thing that comes to mind is that the effect of recency on prominence is mitigated by hierarchical discourse structure induced by subordinating coherence relations, as captured by the Right Frontier Constraint. The chapter develops a theoretical analysis of ARCs as discourse units addressing their own QUDs and connected to their main clauses by coherence relations. The effect of the linear position of an ARC in the sentence (50)–(51) follows automatically from applying the Right Frontier Constraint. In addition, it follows that the at-issue status and challengeability of a sentence-final ARC should depend on whether it is connected to its main clause by a subordinating or a coordinating relation. Most of the time, ARCs are discourse-structurally subordinate to their main clauses. For instance, in (51) the ARC is an *Explanation* of the main clause. ARCs that are connected to their main clauses by coordinating coherence relations are known in the literature as *continuative* relative clauses, as in the following examples from Holler (2008). In (52) the relation between the clauses is *Narration*, whereas in (53) it is *Contrast*.

(52) Oskar traf einen Bauern, den er *dann* nach dem Weg fragte.
Oskar met a farmer whom he then for the way asked
'Oskar met a farmer, whom he then asked the way.'

(53) Oskar machte einen Versuch, der *aber* restlos scheiterte.
Oskar made an attempt which however completely failed
'Oskar made an attempt, which however completely failed.'

The Right Frontier Constraint predicts that in the case of subordination in (51) a rejection like *No, he didn't!* should be equally able to target the main and the relative clause. In (52) and (53), however, the coordinating relation between the main clause and the ARC makes the main clause inaccessible for attachment of new discourse material, and for rejection in particular. Therefore, a rejection should be more likely to target the relative clause. For instance, in (52) a rejection like *No, he didn't!* is more likely to be interpreted as 'No, he did not ask the way', than as 'No, he didn't meet a farmer'. In other words, continuative ARCs even turn out to be 'more at-issue' than main clauses. This prediction has been in the meantime confirmed experimentally by Živković (2016). Where the relation between the main and the relative clause was subordinating, the participants chose a rejection targeting the ARC about half of the time. If the relation was coordinating, a rejection targeting the ARC was preferred in about 75% of the cases.

Živković's study corroborates the thesis proposed in chapter 10 that at-issueness is essentially a matter of prominence, at least as far as the way it manifests itself in the accessibility for rejection is concerned. This result provides further support for the highly context-sensitive, pragmatic nature of (not-)at-issue status and is hard to reconcile with the practice of encoding not-at-issue status of specific words or constructions in the lexicon or the grammar.

Moreover, the reformulation of the Right Frontier Constraint and the discourse-structural subordination/coordination distinction in terms of QUD stacks sketched out in chapter 8 provides the theoretical foundation for the link between prominence and at-issueness. According to Simons et al.'s (2011) definition, a proposition is at-issue if it is relevant to the current QUD. The current QUD is the one on top of the QUD stack. The material associated with the QUD on top of the stack is most prominent and is at issue. It has to be kept in memory as long as the QUD remains open. Once the QUD is handled, closed, and popped off the stack, the corresponding material becomes less prominent and is not at issue any more. In other words, the study in chapter 10 and the subsequent study by Živković (2016) provide both empirical and theoretical support for the pragmatic approach to peripheral components of sentence content.

To conclude, the two studies presented in chapters 9 and 10 demonstrate how the QUD-based notion of discourse structure developed in the first seven papers in this collection can be put to use in the discussion of strongly debated issues in pragmatics. QUD-based discourse structure is needed for a Gricean account for discourse-level

implicatures, a phenomenon that shows that Gricean quantity implicatures cannot be entirely replaced by localist exhaustivity operators in the sentence structure. Furthermore, QUD-based discourse structure establishes a link between at-issue status and prominence, reinforcing the pragmatic approach to at-issueness and confronting the conventionalists with new hard empirical problems. All in all, this shows that the development of the QUD-based theory of discourse structure brings progress to the field of pragmatics.

1.7 Conclusion

In this volume I have collected nine papers that are united by a single research programme. This research programme takes its origin in the general dissatisfaction with the conceptual motivation of the existing sets of coherence relations that constitute the basis of the relational approach to discourse structure. Instead, I have taken up the challenge to develop a goal-based theory of discourse. Although the idea of goal-based discourse structure is quite old, it is still necessary to fill it with concrete content. That means, first of all, to provide a goal-based model of relational inferences—the 'extra content' behind *Elaboration*, *Explanation*, *Contrast*, etc. that we perceive to be present when we interpret connected discourse. I use the notion of Question under Discussion (QUD) as an operationalization of the notion of discourse goal, and develop QUD-based accounts of the inference of subordinating relations such as *Elaboration*, *Explanation* and *Background*, in the absence of explicit markers signalling such relations. I argue that coordinating relations like *Contrast*, *Parallel* and *Narration* require at least some minimal linguistic marking and develop a QUD-based semantics for discourse connectives that encode these relations. I also provide an account of *Correction* as a 'hybrid' coherence relation that is inferred like *Elaboration* but signalled like *Contrast*.

The second task on the way to a fully fledged goal-based theory of discourse is to provide a model of context for the study of context-sensitive linguistic expressions and pragmatic inferences. An important source of motivation for a theory of discourse is how well it can help explain those phenomena and resolve theoretical issues in sentence semantics and pragmatics. The work presented in this volume shows that QUD-based discourse structure is needed for Gricean reasoning towards conversational quantity implicatures, in particular when such implicatures are drawn above the sentence level, and for explaining the variable at-issue status of non-restrictive relative clauses. These findings make an original contribution to the current debate on these central issues in pragmatics.

Even though many more relevant phenomena still await a QUD-based explanation, this is already quite a representative collection that demonstrates the potential of the approach. Possible directions for further development include an application to narrative discourse and causality (coherence relations like *Narration* and *Result*), working out the details of the QUD-based account of discourse prominence and anaphora, and further experimental investigations of the new hypotheses resulting from the proposed

theories.

Bibliography

Amaral, P., Roberts, C., Smith, E. A., 2007. Review of The Logic of Conventional Implicatures by Chris Potts. Linguistics and Philosophy 30 (6), 707–749.

AnderBois, S., Brasoveanu, A., Henderson, R., 2011. Crossing the appositive/at-issue meaning boundary. In: Proceedings of Semantics and Linguistic Theory, SALT'20. pp. 328–346.

Asher, N., 2013. Implicatures and discourse structure. Lingua 132, 13–28.

Asher, N., Hardt, D., Busquets, J., 1997. Discourse parallelism, scope, and ellipsis. In: Peters, S., de Swart, H. (Eds.), Proceedings of SALT 7. Stanford.

Asher, N., Lascarides, A., 1995. Lexical disambiguation in a discourse context. Journal of Semantics 12 (1), 69.

Asher, N., Lascarides, A., 1998a. Bridging. Journal of Semantics 15 (1), 83.

Asher, N., Lascarides, A., 1998b. The semantics and pragmatics of presupposition. Journal of Semantics 15 (3), 239.

Asher, N., Lascarides, A., 2003. Logics of Conversation. Studies in Natural Language Processing. Cambridge University Press.

Asher, N., Vieu, L., 2005. Subordinating and coordinating discourse relations. Lingua 115, 591–610.

Barwise, J., Cooper, R., 1981. Generalized quantifiers and natural language. Linguistics and Philosophy 4 (2), 159–219.

Bäuerle, C. S. R., Schwarze, C., von Stechow, A. (Eds.), 1983. Meaning, Use and the Interpretation of Language. De Gruyter, Berlin.

Beaver, D. I., Clark, B. Z., 2008. Sense and Sensitivity: How Focus Determines Meaning. Blackwell.

Benz, A., Jasinskaja, K., 2017. Questions under discussion: From sentence to discourse. Discourse Processes 54 (3), 177–186.

Benz, A., Kühnlein, P., Sidner, C. (Eds.), 2010. Constraints in Discourse 2. John Benjamins, Amsterdam.

Bott, O., Solstad, T., 2014. From verbs to discourse: A novel account of implicit causality. In: Hemforth, B., Mertins, B., Fabricius-Hansen, C. (Eds.), Psycholinguistic approaches to meaning and understanding across languages. Springer, pp. 213–251.

Büring, D., 2003. On D-trees, beans and B-accents. Linguistics and Philosophy 26, 511–545.

Chemla, E., Spector, B., 2011. Experimental evidence for embedded scalar implicatures. Journal of Semantics.

Chierchia, G., 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In: Belletti, A. (Ed.), Structures and Beyond, Volume 3: The Cartography of Syntactic Structures. OUP USA, pp. 39–103.

Chierchia, G., 2006. Broaden your views: Implicatures of domain widening and the "logicality" of language. Linguistic Inquiry 37 (4), 535–590.

Clark, H. H., 1996. Using Language. Cambridge University Press.

Danlos, L., 2001. Event coreference in causal discourses. In: Bouillon, P., Busa, F. (Eds.), The Language of Word Meaning. Cambridge University Press, pp. 216–241.

Dowty, D. R., 1979. Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and in Montague's PTQ. D. Reidel Publishing Company.

Fox, D., 2007. Free choice and the theory of scalar implicatures. In: Sauerland, U., Stateva, P. (Eds.), Presupposition and implicature in compositional semantics. Palgrave Macmillan, pp. 71–120.

Fox, D., Katzir, R., 2011. On the characterization of alternatives. Natural Language Semantics 19, 87–107.

Gajewski, J., Sharvit, Y., 2009. In defense of the grammatical approach to local implicatures. Natural Language Semantics, 1–27.

Geurts, B., 2007. Implicature as a discourse phenomenon. In: Puig-Waldmüller, E. (Ed.), Proceedings of Sinn und Bedeutung 11. Universitat Pompeu Fabra, Barcelona, Spain, pp. 261–275.

Geurts, B., 2010. Quantity Implicatures. Cambridge University Press.

Ginzburg, J., 1995. Resolving questions, I. Linguistics and Philosophy 18, 459–527.

Ginzburg, J., 1996. Dynamics and the semantics of dialogue. In: Seligman, J., Westerståhl, D. (Eds.), Logic, Language and Computation. Vol. 1. CSLI Publications, Stanford, CA.

Ginzburg, J., 2012. The interactive stance. OUP Oxford.

Ginzburg, J., Fernández, R., Schlangen, D., 2007. Unifying self- and other-repair. In: Arstein, R., Vieu, L. (Eds.), Decalog 2007: Proceedings of the 11th Workshop on the Semantics and Pragmatics of Dialogue. pp. 57–63.

Grice, H. P., 1975. Logic and conversation. In: Cole, P., Morgan, J. L. (Eds.), Syntax and Semantics 3: Speech Acts. Academic Press, New York, pp. 41–58.

Groenendijk, J., Roelofsen, F., 2009. Inquisitive semantics and pragmatics. In: Proceedings of the ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric. University of the Basque Country Press.

Groenendijk, J., Stokhof, M., 1984. Studies in the semantics of questions and the pragmatics of answers. Ph.D. thesis, University of Amsterdam.

Groenendijk, J., Stokhof, M., 2011. Questions. In: van Benthem, J., ter Meulen, A. (Eds.), Handbook of Logic and Language, 2nd Edition. Elsevier, London, pp. 1059–1131.

Grosz, B. J., Sidner, C. L., 1986. Attention, intentions and the structure of discourse. Computational Linguistics 12 (3), 175–204.

Hamblin, C. L., 1973. Questions in Montague English. Foundations of Language 10, 41–53.

Hobbs, J. R., 1985. On the coherence and structure of discourse. Tech. Rep. CSLI-85-37, Center for the Study of Language and Information, Stanford University.

Holler, A., 2008. A discourse-relational approach to continuation. In: Benz, A., Kühnlein, P. (Eds.), Constraints in Discourse. John Benjamins, pp. 249–266.

Hovy, E. H., Maier, E., 1995. Parsimonious or profligate: How many and which discourse structure relations. Tech. rep., University of Southern California.

Hunter, J., Abrusán, M., forthcoming. Rhetorical relations and QUDs. In: Proceedings of LENLS.

Jacobs, J., 1982. Syntax und Semantik der Negation im Deutschen. Fink, München.

Jakobson, R., 1971. Quest for the essence of language. In: Selected writings, Vol. 2: Word and language. Mouton, The Hague, Netherlands, pp. 345–359, first published in 1965.

Jasinskaja, K., 2010a. Corrective contrast in Russian, in contrast. Oslo Studies in Language 2 (2), 433–466.

Jasinskaja, K., 2010b. Modelling discourse relations by topics and implicatures: The elaboration default. In: Benz et al. (2010), pp. 61–79.

Jasinskaja, K., 2012. Correction by adversative and additive markers. Lingua 122, 1899–1918.

Jasinskaja, K., 2013. Corrective elaboration. Lingua 132, 51–66.

Jasinskaja, K., in revision. Not at issue any more, ms. University of Cologne. Submitted to *Journal of Semantics*.

Jasinskaja, K., Karagjosova, E., submitted. Rhetorical relations, submitted to: Matthewson, Lisa & Cécile Meier & Hotze Rullmann & Thomas Ede Zimmermann (eds.): *The Companion to Semantics*. Oxford: Wiley.

Jasinskaja, K., Salfner, F., Freitag, C., 2017. Discourse-level implicature: A case for QUD. Discourse Processes 54 (3), 239–258.

Jasinskaja, K., Zeevat, H., 2008. Explaining additive, adversative and contrast marking in Russian and English. Revue de Sémantique et Pragmatique 24, 65–91.

Jasinskaja, K., Zeevat, H., 2009. Explaining conjunction systems: Russian, English, German. In: Riester, A., Solstad, T. (Eds.), Proceedings of Sinn und Bedeutung 13. Vol. 5 of SinSpeC. Working Papers of the SFB 732. University of Stuttgart, pp. 231–245.

Karttunen, L., 1977. Syntax and semantics of questions. Linguistics and Philosophy 1, 3–44.

Katzir, R., 2007. Structurally-defined alternatives. Linguistics and Philosophy 30 (6), 669–690.

Kehler, A., 2002. Coherence, Reference, and the Theory of Grammar. CSLI Publications.

Kehler, A., Kertz, L., Rohde, H., Elman, J. L., 2008. Coherence and coreference revisited. Journal of Semantics 25 (1), 1–44.

Klein, W., von Stutterheim, C., 1987. Quaestio und die referentielle Bewegung in Erzählungen. Linguistische Berichte 109, 163–185.

Knott, A., 1996. A data-driven methodology for motivating a set of coherence relations. Ph.D. thesis, The University of Edinburgh: College of Science and Engineering: The School of Informatics.

Knott, A., Dale, R., 1994. Using linguistic phenomena to motivate a set of coherence relations. Discourse Processes 18 (1), 35–62.

Knott, A., Sanders, T., 1998. The classification of coherence relations and their linguistic markers: An exploration of two languages. Journal of Pragmatics 30, 135–175.

Koev, T. K., 2013. Apposition and the structure of discourse. Ph.D. thesis, Rutgers University-Graduate School-New Brunswick.

Kratzer, A., 1981. The notional category of modality. In: Eikmeyer, H.-J., Rieser, H. (Eds.), Words, Worlds, and Contexts, New Approaches to Word Semantics. Walter de Gruyter, Berlin, p. 38–74.

Kratzer, A., Shimoyama, J., 2002. Indeterminate pronouns: The view from Japanese. In: 3rd Tokyo conference on psycholinguistics.

Krifka, M., 2001. For a structured meaning account of questions and answers. In: Fery, C., Sternefeld, W. (Eds.), Audiatur Vox Sapientia. A Festschrift for Arnim von Stechow. Vol. 52 of studia grammatica. Akademie Verlag, Berlin, pp. 287–319.

Lascarides, A., Asher, N., 1993. Temporal interpretation, discourse relations and commonsense entailment. Linguistics and Philosophy 16, 437–493.

Link, G., 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In: Bäuerle et al. (1983), pp. 303–323.

Litman, D. J., Allen, J., 1990. Discourse processing and commonsense plans. In: Cohen, P. R., Morgan, J., Pollack, M. E. (Eds.), Intentions in Communication. MIT Press, pp. 365–388.

Lohnstein, H., 2016. Verum focus. In: Féry, C., Ishihara, S. (Eds.), The Oxford Handbook of Information Structure. Oxford University Press.

Malchukov, A. L., 2004. Towards a semantic typology of adversative and contrast marking. Journal of Semantics 21, 177–198.

Mann, W. C., Thompson, S., 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. Text 8 (3), 243–281.

Mauri, C., 2008. Coordination Relations in the Languages of Europe and Beyond. Mouton de Gruyter.

Montague, R., 1970. English as a formal language. In: Visentini, B. (Ed.), Linguaggi nella societa e nella tecnica. Edizioni di Communita, pp. 188–221.

Montague, R., 1973. The proper treatment of quantification in ordinary English. In: Hintikka, K. J. J., Moravcsik, J. M. E., Suppes, P. (Eds.), Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics. Springer Netherlands, Dordrecht, pp. 221–242.

Onea, E., 2016. Potential Questions at the Semantics-Pragmatics Interface. Brill.
Partee, B., Rooth, M., 1983. Generalized conjunction and type ambiguity. In: Bäuerle et al. (1983), pp. 361–393.
Polanyi, L., 1988. A formal model of the structure of discourse. Journal of Pragmatics 12, 601–638.
Potts, C., 2005. The Logic of Conventional Implicatures. Oxford University Press.
Roberts, C., 1996. Information structure in discourse: Towards an integrated formal theory of pragmatics. OSU Working Papers in Linguistics 49, 91–136.
Rodríguez, K. J., Schlangen, D., 2004. Form, intonation and function of clarification requests in German task oriented spoken dialogues. In: Proceedings of Catalog '04 (The 8th Workshop on the Semantics and Pragmatics of Dialogue, SemDial04). pp. 101–108.
Rooth, M. E., 1992. A theory of focus interpretation. Natural Language Semantics 1, 75–116.
Russell, B., 2006. Against grammatical computation of scalar implicatures. Journal of Semantics 23 (4), 361–382.
Sauerland, U., 2004. Scalar implicatures in complex sentences. Linguistics and philosophy 27 (3), 367–391.
Sauerland, U., 2012. The computation of scalar implicatures: Pragmatic, lexical or grammatical? Language and Linguistics Compass 6 (1), 36–49.
Schlenker, P., 2013. Supplements without bidimensionalism, unpublished manuscript. Institut Jean-Nicod and NYU.
Schulz, K., van Rooij, R., 2006. Pragmatic meaning and non-monotonic reasoning: The case of exhaustive interpretation. Linguistics and Philosophy 29 (2), 205–250.
Simons, M., Beaver, D., Roberts, C., Tonhauser, J., 2017. The Best Question: Explaining the projection behavior of factives. Discourse Processes 54 (3), 187–206.
Simons, M., Tonhauser, J., Beaver, D., Roberts, C., 2011. What projects and why. In: Proceedings of Semantics and Linguistic Theory, SALT'20. pp. 309–327.
Syrett, K., Koev, T., 2015. Experimental evidence for the truth conditional contribution and shifting information status of appositives. Journal of Semantics 32, 525–577.
Traum, D. R., 1994. A computational theory of grounding in natural language conversation. Ph.D. thesis, University of Rochester.
Txurruka, I. G., 2003. The natural language conjunction *and*. Linguistics and Philosophy 26, 255–285.
van Kuppevelt, J., 1995. Discourse structure, topicality and questioning. Journal of Linguistics 31, 109–147.
van Rooij, R., Schulz, K., 2004. Exhaustive interpretation of complex sentences. Journal of Logic, Language and InformationIn press.
van Rooy, R., 2003. Questioning to resolve decision problems. Linguistics and Philosophy 26, 727–763.
Živković, E., 2016. The at-issue status of appositive relative clauses: Evidence for a discourse-based approach. Master's thesis, Pompeu Fabra University, Barcelona.

Zeevat, H., 2004. Contrastors. Journal of Semantics 21, 95–112.
Zeevat, H., 2009. Optimal interpretation as an alternative to Gricean pragmatics. Oslo Studies in Language 1 (1), 191–216.
Zeevat, H., 2010. Optimal interpretation for rhetorical relations. In: Benz et al. (2010), pp. 35–59.
Zeevat, H., Jasinskaja, K., 2007. *And* as an additive particle. In: Aurnague, M., Korta, K., Larrazabal, J. M. (Eds.), Language, Representation and Reasoning. Memorial volume to Isabel Gómez Txurruka. University of the Basque Country Press, pp. 315–340.