Description
Grammar
for Discourse

Description Grammar for Discourse

een wetenschappelijke proeve op het gebied
van de Filosofie

Proefschrift

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de Rector Magnificus
prof. mr. S.C.J.J. Kortmann
volgens besluit van het College van Decanen
in het openbaar te verdedigen op
woensdag 5 december 2007
om 10.30 uur precies
door

Nora van Leusen

geboren op 15 oktober 1967 te Arnhem

Promotores:
Prof. dr. R.A. van der Sandt
Prof. dr. ir. R.J.H. Scha (Universiteit van Amsterdam)

Manuscriptcommissie:
Prof. dr. M.V.P. Slors
Prof. dr. N.M. Asher (University of Texas, Austin)
Dr. C. Gardent (CNRS Director of research, LORIA/Nancy)

© 2007 Noor van Leusen
ISBN 978-90-9022305-6
Typeset with LaTeX 2ε
Cover design by Marianne Stoop
Printed by Gildeprint Drukkerijen, Enschede

Voor Reinhard

Acknowledgements

The research for this work started at the University of Amsterdam and, after interludes in Saarbrücken and Stuttgart, resulted in a dissertation at the Radboud University in Nijmegen.

As an effect of these travels, I am now blessed with two promotors, Rob van der Sandt and Remko Scha. I would like to thank Remko, the most friendly and patient person alive, for sharing his views on discourse semantics with me in most enjoyable conversations. To Rob I am deeply grateful for instigating me to take up my research again and finish the book. Thank you for many hours spent on discussing general issues in semantics and presupposition theory in particular, for being *very* clear about planning and time-constraints, for occasional trips to the Ooijpolder and to musical events, and for instructions on how to pickle salmon. You've been wonderful.

I would also like to thank Henk Zeevat for inspiring discussions during my years in Amsterdam, and my colleagues there, in particular the one-time inhabitants of room 337 of the PC Hoofthuis, for making me feel at home at Alfa-Informatica. During a later period at the ILLC, Paul Dekker has been extremely supportive and helpful.

In recent years, I joined the small but sociable philosophy of language group in Nijmegen led by Rob van der Sandt. In addition to Rob, I would like to thank Bart Geurts, Emar Maier, Janneke Huitink, and Corien Bary for taking me up in their midst. I much enjoyed your company, and the weekly gossiping during 'vakgroepslunch'. You're a talented bunch of people and I hope you will all find your way in Academia as succesfully as you deserve.

Contents

1 Introduction **11**

2 Logical Description Grammar **21**
2.1 Introduction 21
2.2 Tree Adjoining Grammar 22
2.3 Description Grammars 26
2.4 LDG in Detail 30
2.4.1 Overall Working of the Grammar 31
2.4.2 Parsing an Ambiguous Sentence: An Example 33
2.4.3 General Descriptions 36
2.4.4 Positive and Negative Anchoring 37
2.4.5 Input Descriptions 40
2.4.6 Lexical Descriptions 41
2.5 Conclusion 44

3 LDG for Discourse Analysis **47**
3.1 Introduction 47
3.2 A Framework of Discourse Interpretation 49
3.3 Sentence and Discourse Theory Integrated 53
3.4 Discourse Structure and Discourse Relations 55
3.5 Incremental LDG 64
3.5.1 Circumscription of Lexical Nodes 65
3.5.2 Input Descriptions 66
3.5.3 The Incrementation Operation + ... 68
3.5.4 Reasoning about Discourse Descriptions 70
3.5.5 Attachment at the Right Frontier 77
3.5.6 Deterministic Discourse Processing 81
3.6 Summary and Loose Ends 84

4 Discourse Semantics **89**
4.1 Introduction . . . 89
4.2 Discourse Meaning . . . 91
4.3 Nonmonotonic Discourse Contributions . . . 96
4.4 Context Sensitive Compositional Semantics . . . 98
4.4.1 Finegrained CDRT: Basic Concepts and Notation . . . 100
4.4.2 Decorating Tree Structures with Local Contexts . . . 103
4.4.3 Lexical Descriptions of some Basic Word Classes . . . 106
4.4.4 Discourse Relations, Connectives, and Adverbials . . . 110
4.4.5 Global Consistency and Informativity Conditions . . . 124
4.4.6 Inferring Semantics from Discourse Descriptions . . . 131
4.4.7 Testing Context Sensitive Constraints . . . 135
4.5 Conclusion . . . 143

5 Anaphoricity **147**
5.1 Introduction . . . 147
5.2 Properness and Pronoun Resolution . . . 149
5.2.1 Inferring Binding Relations . . . 150
5.2.2 Semantic Accessibility and Parallelism . . . 153
5.2.3 Kataphors . . . 155
5.2.4 Summing up . . . 156
5.3 Syntactic Binding Principles in LDG . . . 157
5.4 Discourse Topic Structure and Salience . . . 160
5.5 Accommodating Binders . . . 169
5.6 Conclusion . . . 172

6 Presupposition Theory **175**
6.1 Introduction . . . 175
6.2 Binding and Satisfaction in Tandem . . . 178
6.2.1 Presupposition Triggers in the Lexicon . . . 179
6.2.2 Reasoning about Presuppositional Conditions . . . 184
6.2.3 Binders can be Globally Accommodated . . . 187
6.2.4 Referential and Existential Readings of Definites . . . 190
6.2.5 Accommodating Weak and Strong Readings . . . 191
6.2.6 Partial Matching and Suitability . . . 193
6.2.7 Bridging . . . 194
6.3 Consistency and Informativity Revisited . . . 198
6.3.1 Global Consistency and Global Informativity . . . 198
6.3.2 Van der Sandt's Acceptability Conditions . . . 202
6.3.3 Conclusion - Ways to Go . . . 208
6.4 Generalised Accommodation . . . 210
6.4.1 Some Examples of Nonpresupposing Discourse . . . 212

6.4.2 Content Modification—
Local, Intermediate, and in the Main DRS 215
6.4.3 Preferences 220
6.5 The Extended Theory Applied 226
6.5.1 Binding Relative to an Overt Antecedent 227
6.5.2 Nonpresupposing Conditional 229
6.5.3 Presupposing Conditional—
Introducing Local Justification Conditions 231
6.5.4 Trapping through Direct Binding 233
6.5.5 Intermediate Readings 235
6.5.6 Bridging 237
6.5.7 Nonpresupposing Disjunction 241
6.5.8 Denial 244
6.6 Conclusion 248

7 Conclusion **253**

A A Fine-grained Compositional DRT **259**

B Axioms of the Grammar **269**

Bibliography **277**

Summary **291**

Curriculum Vitae **295**

Chapter 1

Introduction

Anyone can tell that conversations and texts have meanings. They come with some content which is conveyed from the speaker or writer to the hearer or reader of the discourse. Moreover, it does not take long to see that texts and conversations are built up from smaller units, such as sentences, phrases, and words, and that there is some sort of systematic relation between the composition of a text or conversation from smaller units and its meaning. Clearly, the text in (1.1a) does not mean the same thing as the one in (1.1b).

(1.1) a. *Jack opened the door. He turned on the light.*

b. *Jack turned on the light. He opened the door.*

The order of the sentences matters. From (1.1a) we conclude that Jack turned on the light *after* he opened the door. Furthermore, we tend to assume that the light he turns on is a light in the room he just opened a door of. This presupposes that before Jack opened the door and turned on the light, the light was off. In fact, in order to turn on the light, he must partly or completely have entered the room, so it is likely that just after he turned on the light he is in the room. From (1.1b) we conclude that Jack turned on the light *before* he opened the door. The light he turns on must be a light in the corridor or location outside the room of which he opens a door. This light must have been off, before Jack turned it on. It is not clear whether Jack enters the room. If we think of the discourse in (1.1b) as a short narrative, or a description of what happened, that description ends just after Jack opens the door.

The example illustrates, in a nutshell, what are the main concerns of *discourse theory*. If a language user reads a text, or takes part in a conversation, he constructs a meaning or interpretation of it in his mind. Conversely, if he wants to write a text, or contribute to a conversation, he translates what he has in mind into words and sentences. Discourse theory investigates the systematic relation between linguistic form and meaning, in particular where the construction and interpretation of larger discourse units is concerned.

It tries to explain how language users handle implicitness in discourse, that is, how they fill in what is not explictly said and obtain cohesive interpretations. It offers an analysis of the context dependence of sentences in their discourse context. It addresses, for instance, the particular contribution to discourse meaning of words and phrases that link up to events or individuals previously mentioned in the discourse. These are called *anaphors*. The pronoun 'he' in (1.1a-b) is an examples of this. Definite noun phrases, such as 'the light' and 'the door', are sometimes called anaphors as well. They refer to entities which are 'given', or can presumed to be present in the context of interpretation. Because they tell us something about what the speaker or writer presupposes to be the case they are also often categorised as *presupposition triggers*. Sometimes an anaphor has an explicit *antecedent*, that is, a word or phrase that precedes it in the discourse, and which introduces the entity it refers to. For example, in (1.1a-b) 'Jack' is the antecedent of 'he'.

Discourse grammar is a branch of discourse theory which most rigorously implements the idea that discourse is composed of clauses and larger discourse units just like sentences are composed of words and phrases. Discourse, like sentences, has syntactic structure, and the construction of discourse is rule-based. While the meaning of a sentence is composed of the meanings of the words it consists of, the meaning of a discourse is composed of the meanings of sentences and so-called discourse relations. In (1.1), for example, it is assumed that each discourse is composed of two sentences that are linked by a *narration* relation. Although a discourse relation is not expressed by a visible lexeme, it has a syntactic and semantic contribution of its own, just like a word. The main semantic contribution of a narration relation is to establish a temporal relation between the events described by the two sentences in the discourse. The event described by the first sentence precedes the event expressed by the second sentence in time.

In this work we present a framework of discourse interpretation that includes a discourse grammar. The purpose of our enterprise is to put forward what might be called a common language for the specification of discourse theory. Much of current research in linguistics and discourse interpretation focuses on highly specialised issues. Moreover, it employs a number of different formal representation systems, so that the analyses provided are sometimes hard to compare. Though it is unavoidable that alternative treatments of one and the same topic arise in a field of research where we meet with uncertainties of various kinds, such a dispersion of our findings need not be accepted. Even when we concentrate on details or isolated phenomena, linguistic theory is essentially a theory of the *whole* of a language user's linguistic capabilities. Analyses of different topics should be combined so that they may help to evaluate and strengthen each other.

We think it worthwile to strive for unity in formal treatment, and to develop a shared discourse theoretical perspective. The framework of discourse interpretation that we put forward is a flexible system which allows much freedom in setting the parameters of discourse theory. We hope that it may help to integrate

and compare findings that are now spread over different approaches or spelled out in different formal theories. To show the possibilities of the framework it is applied in the domain of anaphoricity and presupposition theory in this work.

The framework is a *reasoning system* which includes a discourse grammar. The grammar represents the syntactic and semantic 'rules' in the mind of a language user. It models his linguistic knowledge. As may be gathered from our brief discussion of (1.1), however, linguistic knowledge on its own is not enough to obtain the meaning of a discourse. A language user employs his world knowledge and his knowledge of what is normally the case as well. He is able to infer a most plausible interpretation of a discourse, given his linguistic and his nonlinguistic knowledge. Therefore, our framework of discourse interpretation also includes a reservoir of world knowledge, and a so-called default reasoning component and preference system.

Another salient characteristic of the discourse model is that the grammar is a Logical Description Grammar. A standard discourse grammar generates linguistic representations or discourse representations for a given discourse. Usually, linguistic representations are *tree structures*,[1] decorated with syntactic labels and features, and semantic values. If a single sentence or discourse has more than one possible analysis or interpretation, the grammar generates a different representation for each of them. We say, in that case, that the discourse is *ambiguous* or *underspecified*. The following sentence is a example of what is usually called structural ambiguity.

(1.3) *Alan filmed a cat sitting on the roof.*

Depending on how we link the phrase 'sitting on the roof' to the other parts of the sentence we get different meanings. If it modifies the object 'a cat' it indicates a property of the animal: the cat that Alan filmed was sitting on the roof. On the other hand, if it attaches to the verb phrase 'filmed a cat' it specifies the manner in which Alan filmed a cat, namely while he (Alan) was sitting on the roof. Another example is given in (1.4); this is a common case of underspecification.

[1] A tree structure is a graph with certain special characteristics. A *graph* is a finite set of nodes with a relation on it. A *tree structure* is a graph with a single root, or topnode. If there is more than one node, the root is connected to one or more 'daughter nodes', which themselves can be the topnode of a subtree. Subtrees are disjunct. Here is an example; the nodes are indicated by letters.

(1.2)

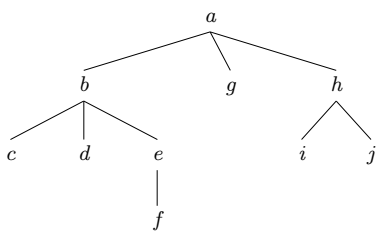

(1.4) *John fell. Bill pushed him.*

We may fill in the implicit relation between the two events differently. Either John fell, and subsequently, Bill pushed him. Or we infer that John fell because Bill pushed him, so the pushing took place before the falling.

Ambiguity and underspecification may come in various forms and guises in a given language, and discourse abounds with it. A grammar formalism can capture the interpretation of an ambiguous sentence or discourse by producing all possible linguistic representations and corresponding meanings. However, it can also produce a single representation, which underspecifies the syntactic or semantic properties that are variable and differ across the possible representations. *Description grammars* are of the latter kind.[2] For any given sentence or discourse they generate a single description, which specifies all that is invariable about the input, and leaves open the options when there is a choice. In the case of (1.3), the description will say that 'sitting on the roof' is a constituent within the sentence but it will not determine its exact syntactic position. In (1.4), it will say that there is a relation between the events expressed by the two sentences, but it will not specify which one.

In a description grammar there are two levels of linguistic analysis, namely that of the description, which can be partial, and that of the linguistic structures or representations it describes. For any given input, the possible representations can be obtained by reasoning about the tree structures that 'fit' the description and the semantic content that may possibly fill the gaps. In particular, we will speak of a Logical Description Grammar or LDG. This serves to emphasize that the grammar formalism in our framework is in fact a logical theory, that is, a set of statements, or *constraints* in a logical language. Sentence and discourse descriptions are likewise sets of statements in a logical language. The combined theory of the grammar and a discourse description describes the syntax and semantics of a given discourse. The linguistic representations of a discourse can be obtained by reasoning about the models that verify the combined theory.

There are various reasons why the use of descriptions in linguistics is attractive. One is the very simple one of having compact representations of what must otherwise be summed up in a list. Another is the computational advantage of not having to backtrack too soon in practical implementations of natural language processing. From a theoretical point of view, an important reason is that with partial descriptions we can state most directly and precisely what we know about the words and sentences they represent, and what can assumed to be their contribution to the syntax and semantics of the larger whole of a discourse. We can abstract over what is variable in their behaviour. This, in principle, leads to conceptual accurateness and transparancy in our grammar and discourse descriptions. Finally, since the description method applies not just to individual words

[2]See for instance Vijay-Shanker (1992), Rogers and Vijay-Shanker (1992), Muskens (1995), Gardent and Webber (1998), and Kallmeyer (1999).

or particular discourses, but to linguistic theory as such, there is a great flexibility and generality in the approach.

It is the last point we take as a lead in our discourse model. While the components of the framework we have on offer are not different from what is present in other approaches, we present a meta-theoretical perspective which is new. The description method allows us to talk about discourse theory as such, and to consider different ways of implementing it. The framework may be seen as a test-bed for discourse theory, and a tool for comparing and integrating different approaches towards the same linguistic topic. We think this is particularly useful in a field like discourse semantics, where, we would claim, the theories that are developed are open ended in two ways. First, it is not very clear how to delineate the object of research and determine what a discourse theory should account for. Second, there is the issue of what power or complexity the theory itself might or might not have.

As for the first point, we will try to stay on safe ground by concentrating on the observable linguistic facts, that is, words, sentences, and texts in English, and the possible readings that sentences can get in a given discourse context, according to our own intuitive judgement and that of other language users. The latter is already slippery ground, though. Beyond this is the vista of the communication process as whole, and all the aspects of human language behaviour. We will stay away from that, but it cannot be denied that the things that people say or write are part of their communicative behaviour. We will sometimes refer to rules of communication and language users' purposes, in order to explain our 'observable facts'. Furthermore, as is common nowadays, we adopt a cognitive linguistic perspective. We will talk about the language user's grammatical knowledge, and will consider his language capacity as a cognitive capability.

As for the technical properties of the formalism, we allow ourselves a discourse theory that is highly powerful. The description grammar is in the class of context-sensitive grammar formalisms. The grammar is a set of hard constraints: a discourse must satisfy these constraints, if not, it is predicted to be ungrammatical. The grammar is complemented with a so-called preference system and a common sense reasoning component. Preferences and common sense inferences generate gradient predictions, and these predictions are cancellable in the presence of new information. The preference system raises the complexity of the discourse theory, but it seems necessary to accept this if one aims at a linguistic coverage that is anywhere near complete. At the same time, it is good practice to explain the phenomena we are looking at in the simplest possible way, so in general we will try to formulate our account in terms of plain grammatical constraints first, and only if that does not work resort to preferences.

Existing discourse theories vary greatly in what they cover or intend to explain, and in the formal tools they employ. They range from those that provide a logic to account for relatively few aspects of the semantics of discourse, for instance

Discourse Representation Theory,[3] File Change Semantics,[4] Dynamic Montague Grammar,[5] Dynamic Predicate Logic,[6] and Update Semantics,[7] to those that have a broad linguistic coverage and are meant to account for the macro-structure of discourse, such as the Linguistic Discourse Model,[8] Rhetorical Structure Theory,[9] Segmented Discourse Representation Theory,[10] and (Lexicalised) Tree Adjoining Grammar for Discourse.[11] They may also build on theories of planning and goal satisfaction, like Centering Theory,[12] or focus on dialogue interpretation.[13]

The description grammar framework is powerful enough to implement the greater part of the linguistic phenomena that these theories cover. Due to the description method, the parameters of the discourse theory can be varied, and the constraints can be set exactly as one wants. Thus, for any given linguistic topic, a particular account may be tested by implementing it, and trying out alternative accounts. On the basis of such comparisons we can hopefully make some motivated choices, not just with respect to the topic under consideration, but also with respect to the discourse theory in which it is embedded.

In chapter 2, 3, and 4 of this book the LDG framework of discourse interpretation is developed. In chapter 5 and 6 the model is applied in the domain of anaphoricity and presupposition theory. In the second of these chapters we take up two competing theories of presupposition interpretation in discourse, and show that they can fruitfully be combined in a single treatment.

Chapter 2 introduces the concept of description grammars and presents the Logical Description Grammar formalism that we will employ. In chapter 3, the LDG formalism is tuned to the treatment of discourse, and the framework of discourse interpretation is put forward. The emphasis in this chapter is on discourse syntax. We discuss the place of discourse relations in the grammar formalism, and their relation with underspecification. We zoom in on the relation between discourse and sentence grammar. Finally, we discuss the ins and outs of the framework as an *incremental* model for discourse interpretation. We assume that language users process discourse step by step. The grammar generates a discourse description for the discourse processed so far, and each newly processed contribution to the discourse induces an incrementation of the discourse description.

[3] Kamp (1981)
[4] Heim (1982)
[5] Groenendijk and Stokhof (1990)
[6] Groenendijk and Stokhof (1991)
[7] Dekker (1993)
[8] (Polanyi 1988; Scha and Polanyi 1988)
[9] Mann and Thompson (1988)
[10] Asher (1993)
[11] Webber et al. (1999)
[12] Grosz and Sidner (1986)
[13] like for instance Hamblin (1971), Carlson (1983), Ginzburg (1994), Walker (1996a), and Hulstijn (2000).

After each incrementation step, the language user may reason about the discourse description, and obtain the possible interpretations of the discourse. Chapter 3 spells out the basic properties of the grammar formalism and technicalities of the reasoning process.

Chapter 4 unfolds the semantics of the formalism. It defines the notion of *discourse meaning* dynamically, as the update of what is implicitly taken for granted by the participants in the discourse with what is being said in it. After each discourse contribution a newly updated discourse meaning can be inferred from the discourse description. Discourse descriptions describe discourse tree structures. Each node in a discourse tree comes with a semantic value, and the semantic value of the topnode is constructed compositionally from the semantic values of the nodes it dominates. The semantic representation language—which is a sublanguage of the description language—is a variant of Compositional Discourse Representation Theory, put forward in Van Leusen and Muskens (2003). Fine-grained CDRT, as we call it, is informally introduced in the chapter and defined in detail in appendix A.

The main issue in chapter 4 is the treatment of context dependence, or context sensitivity in discourse, and how to account for this compositionally. With a view to the analysis of anaphors, presupposition triggers, and context-sensitive words and phrases in general, we assign not only semantic values, but also *local contexts* to the nodes in a discourse tree. Local contexts represent precisely the contextual information that is relevant for the interpretation of the discourse constituent represented by the node to which they are assigned. Anaphors, presupposition triggers, and discourse relations are characterised in the lexicon of the grammar in terms of, among others, conditions on their local contexts. Furthermore, general constraints on the informativity and consistency of discourse constituents in their local contexts can be included in the grammar.

In all of this, the description method and the overall architecture of the framework lead to a smooth integration of grammatical levels, and of sentence and discourse theory. Lexical descriptions describe the syntactic and semantic contribution of words in parallel, and this carries over to discourse descriptions. The grammar is simply a set of constraints which grammatical discourse must satisfy. Syntactic and semantic constraints naturally interact in delimiting grammatical discourse. Although we have not specified phonological and morphological constraints, it seems reasonable to expect that these components can be incorporated in the grammar as well.

As for the place of pragmatics in our framework, matters are somewhat different. Meaning effects that are cancellable, or which result in acceptability judgements that are gradient, are to be explained in terms of the language user's preferences, and we decided to call parameters of the discourse theory that are captured by means of preferences pragmatic. Since the whole of our framework is a reasoning system in which grammatical constraints and preferences interact freely, pragmatics, in our sense of the term, is a fully integrated component of

discourse theory. It is a question for future investigation, however, how much of what is commonly called pragmatics is covered by our notion. It remains to be seen, for instance, exactly how a characterisation of discourse contributions in terms of their appropriate use in communication, and the intentions, commitments and beliefs of the participants could be given in our framework of discourse interpretation.

Chapter 4 completes the presentation of our description grammar framework. The basic parameters of the discourse theory are specified and the general approach to discourse analysis should be clear at this point. The framework is an open work environment, in which some aspects of discourse theory are formally spelled out, and others await further elaboration. It is flexible in the sense that even in the case of basic parameters, which determine, for instance, the type of structures that discourse descriptions can describe, or the semantic representation language, different choices could be made. For easy reference, the basic axioms and definitions of the LDG formalism are summed up in appendix B.

In chapter 5 the analysis of context sensitive elements introduced in chapter 4 is taken up, and a treatment of anaphoricity as an independent parameter in the discourse theory is put forward. Anaphors introduce underspecified discourse markers. Anaphora resolution is an inference task, in which the language user must find out what discourse marker present in the local context of the anaphor potentially binds it, and determine which of the potential binders results in a most preferred reading of the discourse. The inference task is constrained by general semantic conditions on binding, as well as constraints that are particular to the class of anaphors under consideration. For instance, for personal pronouns, we assume that their binders must be entities that are salient, and in a domain of elements determined by the local discourse topic. We touch upon the analysis of so-called bridging anaphors, and discuss the possibility of *accommodating* binders. This means that the presence of a suitable binder in the local context of the pronoun is presumed, even though there is no explicit antecedent in the discourse context which introduces it.

In chapter 6 we zoom in on presupposition theory. We combine the approach to presupposition theory proposed in Karttunen (1974) and others, with the approach put forward in Van der Sandt (1992) and Geurts (1999). The first is sometimes called the *satisfaction* theory of presupposition. It treats presuppositions as tests on the context of interpretation. For a sentence to be felicitous or interpretable, the context of interpretation must entail or satisfy its presuppositions. Sometimes a presupposition can only be satisfied by strengthening the context of interpretation; this is accommodation. The second approach may be called the *anaphoric* theory of presupposition. Its central idea is that presuppositional material must be bound in the discourse context of the trigger. If it cannot be bound, it must be inserted at some semantically accessible location in the discourse context; this is accommodation in the sense of the anaphoric approach.

Starting out from Karttunen (1974), we view presuppositions as conditions on local contexts, and implement the relation of satisfaction as a relation of entailment. On the other hand, we analyse presupposition triggers as anaphors, in line with the anaphoric approach. Anaphoric aspects of presupposition interpretation can then be explained in terms of the interaction with independent constraints on binding in the discourse theory. Presuppositional conditions can only be satisfied given the binding of the trigger.

Fusing the two approaches, we find that each comes with its own notion of accommodation. In the satisfaction approach, this is the strengthening of the initial context of interpretation of a conversation with general background knowledge, as induced by the satisfaction of presuppositions or arbitrary other constraints. In the anaphoric approach, it is the insertion of precisely the presuppositional material in accessible locations in the semantic representation of the discourse. Though these are qualitatively different forms of accommodation, both are the effect of the language user's effort to obtain a possible and most preferred discourse meaning, and we shall argue that if the interpretation task calls for it, they may co-occur. They will be implemented in our discourse semantics accordingly.

Presupposition theory is an integrated part of discourse theory. When a sentence contains a presupposition trigger, the satisfaction of the presuppositional condition is only one factor in its interpretation. The hearer may use the total of his grammatical knowledge to infer the possible grammatical readings of the sentence. Reasoning on the basis of his preferences and world knowledge, he will determine the most preferred readings among these. Thus, what comes out as a most preferred reading is the result of the interaction of all grammatical constraints and preferences involved. This explains that sentences that may be said to evoke weak presuppositions can get strongly presupposing interpretations.

Generally speaking, the predictions of our combined treatment conform to those of the anaphoric approach, if the constraints on informativity and consistency and the preferences in the discourse grammar are set in the manner of the anaphoric account. The explanatory force of the satisfaction approach is preserved in the assumption that presupposition triggers introduce context sensitive conditions. As such, they incite the reasoning process that eventually results in the binding of presuppositional material, either by an explicit antecedent, or by accommodated material. Overall, the two approaches are overlapping, and can be seen to complement, rather than rival each other.

As spin-off of our implementation, we will see that various aspects of presupposition theory that received a procedural treatment in the anaphoric approach to presupposition can be treated declaratively in our underspecification framework. Furthermore, given the flexibility of the account, a unified treatment of the diverse class of presupposition triggers may be provided, while still doing justice to their individual differences. Finally, along the way, we offer an interesting, if not conclusive evaluation of the role of informativity conditions in our discourse theory. Informativity is a typical example of a discourse parameter of which it

is not quite clear whether it belongs to discourse semantics or to pragmatics. We lean towards the latter point of view but the reader is invited to join in the discussion.

We wrap up in chapter 7.

We would like to end with a note on coverage, and the usefulness (or unusefulness) of formal theory in discourse analysis. In any theory of discourse interpretation, there is a trade-off between formal precision and empirical coverage. To put it strongly, only a crumb of what one would like to say about the linguistic data and their interpretation can be specified in a formal theory by spending a reasonably acceptable amount of time on it. Still, the analysis must be as precise and explicit as possible. That is why, in general, linguists use formal theories, or logics, or representation formalisms. At the same time, we want to stay close to the empirical data and give a relatively extensive analysis. In practice we will walk the middle road between formal rigour and broad linguistic coverage by being formally precise about what we consider to be basic parameters of discourse theory, while also leaving much of what the discourse theory as a whole is intended to cover unspecified.

It follows from this that we will sometimes have more to say about the linguistic data than is reflected in the formal theory. The conceptual analysis can be richer than the formal treatment, and there is no sense in which the latter is more important than the former. We may add that to capture the facts that one does understand in a formal theory can be difficult, but to avoid formalising inadvertently what one does not fully understand can be just as hard. This is where plain, informal discussion wins over formalisation and must certainly be given its due.

Chapter 2

Logical Description Grammar

2.1 Introduction

The discourse model which is developed in this dissertation consists of a *Logical Description Grammar* or *LDG* for discourse analysis, a reservoir of world knowledge, and a reasoning component. In this chapter we explain what LDGs are and how they work, and set some pointers to the relevant literature.

Description grammars for sentence analysis have been proposed in Vijay-Shanker (1992), Rogers and Vijay-Shanker (1992), Muskens (1995), Pinkal (1996) and Kallmeyer (1999), among others. Description grammars for discourse analysis have been proposed in Gardent and Webber (1998) and Duchier and Gardent (1999). Whereas most natural language grammars generate tree structures as linguistic representations, LDGs generate *logical descriptions* of tree structures. Since these descriptions can be partial, in the sense that more than one tree may satisfy a given description, aspects of the described tree structure may be underspecified. Due to this feature, LDGs have certain attractive properties.

For one, the derivation or incrementation of descriptions is a monotonic, non-destructive process. That is, no information is removed from a description when it is incremented. This is different in a related formalism such as Tree Adjoining Grammar, where the adjunction operation changes structural relations in the input tree. Furthermore, underspecification or ambiguity occurring in linguistic expressions can be represented in an intuitively plausible manner through underspecified representations. Finally, in description grammar any type of linguistic feature which may be relevant for the analysis of a sentence or discourse can be described in one and the same representation language. Thus, the grammar formalism allows for a smooth integration of constraints coming from different grammatical levels, such as syntax, semantics and—in the case of discourse analysis—pragmatics. It also supports a unified account of sentence and discourse analysis.

The structure of this chapter is as follows.

As a preliminary, in section 2.2 the main features of Tree Adjoining Grammar (TAG) are discussed. Description grammar can be viewed as an extension or modification of TAG. Next, in section 2.3, the basic idea underlying description grammars is explained: description theory can be used to underspecify features of linguistic representations. In section 2.4, the central part of this chapter, we zoom in on the specific variant of description grammar that will be employed in our discourse model in later chapters. This is the LDG formalism proposed in Muskens (2001). The section starts with an introduction to the various levels of analysis of the grammar and discusses the components of the grammar and their working in detail. The chapter is summed up in section 2.5.

2.2 Tree Adjoining Grammar

Logical Description Grammar can be viewed as a variant of Tree Adjoining Grammar (TAG). In TAG, linguistic representations are composed from 'elementary trees' which are provided by the lexicon. Parsing a sentence comes down to putting together the elementary trees triggered by the lexemes occurring in the sentence. This is the same in description grammar, be it that in description grammar there is an additional level of representation, namely the level of descriptions of tree structures.

TAG was first introduced in Joshi et al. (1975), as a grammar formalism whose generative capacity is adequate for the treatment of natural language. As was argued in, among others, Shieber (1985), the generative capacity of context-free grammars is not sufficient for natural language. TAGs are 'mildly context-sensitive', that is, slightly more powerful than context-free grammars. TAG has been applied to various kinds of long distance dependencies in linguistic description. Since TAG was put forward, its claim to adequacy for natural language analysis has been put to the test, and this has led to the formulation of additional constraints and various extensions, one of which is LDG. Description grammars are at least as powerful as TAGs.[1]

For in-depth introductions to TAG, we refer to Joshi et al. (1975), Joshi (1987), Abeillé (1993), or Kallmeyer (1999). Here, we only present some necessary background information for our explanation of description grammar. The discussion is not meant as an overview of the TAG-literature, but it should give the reader an easy access to it.

Whereas standard phrase structure grammars are defined in terms of rewriting rules on strings, TAGs are defined in terms of operations on tree structures. TAGs generate tree structures rather than strings. In a TAG, a syntactic structure for a

[1]Exactly how much more powerful they are depends on the specific type of the grammar formalism under consideration; this is a matter for further investigation, see for instance Kallmeyer (1999).

natural language sentence is derived by 'putting together' elementary trees, each of which carries a lexical item at one of its leaves. For example, the syntactic structure for the sentence 'John sneezed'[2]—

(2.1)

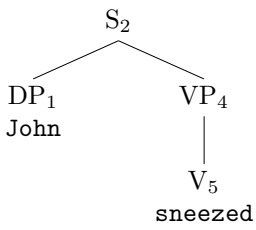

—can be constructed, or derived, from the elementary trees for 'John' and 'sneezed':

(2.2)

(2.3)

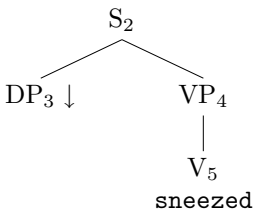

We will say more about the derivation process in a moment. Formally, a TAG consists of a set of *initial* elementary trees, a set of *auxiliary* elementary trees, a set of *terminal* symbols, and a set of *nonterminal* symbols. In the context of natural language processing, the elementary trees represent the syntactic structure and other grammatical features which a lexeme of a particular word class contributes to the syntactic representation of a sentence. Lexemes are terminal symbols and syntactic categories are nonterminal ones. For expository purposes the nodes in the trees in our examples are numbered as well; the numbers do not form part of the elementary trees, however.

Initial elementary trees, such as the two featuring in (2.2) and (2.3), often contain *substitution nodes*. Substitution nodes are nodes on the frontier of the elementary tree labelled by a nonterminal symbol. They are marked by ↓. Linguistically, a substitution node represents a constituent which the lexemes that select the elementary structure subcategorise for. Auxiliary trees—see the example in (2.4) further on—differ from initial trees in that one of the nodes on their frontier, labelled by a nonterminal, is a so-called *foot-node*. This node is annotated with an asterisk (*). The label of the foot-node must be identical to the label of the root node of the elementary tree. Auxiliary trees introduce recursion in a linguistic structure.

[2]The internal syntactic structure of the DP is suppressed in these representations. The same goes for the AP further on.

Derived trees are produced by combining elementary trees, or previously derived trees. The *yield* of a derived tree is the string of terminal and nonterminal symbols we find at its leaves when going through the tree from left to right. The *language* generated by a TAG is the set of terminal strings yielded by the trees that can be derived by the TAG. A TAG is *adequate* for a given natural language if it generates the set of all grammatical expressions of the language and no others, and if the tree structures generated are syntactically correct, given the linguistic theory that is adopted.

Tree structures are derived through the operations 'adjunction' and 'substitution'. For example, the syntactic tree for 'John sneezed' in (2.1) is derived from the two initial trees in (2.2) and (2.3) by substituting the tree headed by the node numbered 1 for the node numbered 3. *Substitution* is the replacement of a substitution node in a tree by another tree. The substitution node and the root of the tree to be substituted must carry the same syntactic label. Substitution is typically used to link a predicate to the arguments or complements it subcategorises for.[3] *Adjunction*[4] is typically used to model the (recursive) application of modifiers. The operation makes use of auxiliary trees. As stated above auxiliary trees contain a unique node which is marked as a *foot-node* and which carries the same nonterminal label as the root. Here's an example:

(2.4)

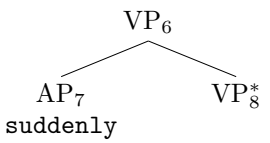

Adjunction is an insertion operation on trees which replaces a subtree rooted at a node called the *adjunction site*. Adjunction replaces the subtree by an auxiliary tree, while attaching the subtree at the foot node of the auxiliary tree. Adjunction may only take place when the labels of the root node of the subtree and the foot node of the auxiliary tree are identical. For example, the syntactic tree for the predicate 'suddenly sneezed' is derived by adjunction of the auxiliary tree for 'suddenly' in (2.4) to the elementary tree for 'sneezed' in (2.3) at node 4, the adjunction site. This results in the structure shown in (2.5). Subsequent substitution of the elementary tree for 'John' in (2.2) at the substitution site in (2.5) will derive a complete syntactic structure for 'John suddenly sneezed'.

[3]However, as is shown in Kroch (1987) and Abeillé (1988a), Abeillé (1988b), this correspondence breaks down for complement clauses when the interaction with long-distance dependencies is considered. The matrix sentence which subcategorises for a complement clause should feature a substitution node, but the interaction with long distance dependencies forces it to be a foot node. In Kallmeyer (1999) it is argued that this is one of the problems of TAGs solved in LDGs.

[4]Adjunction is referred to as 'insertion' in Schilder (1997). It appears in the form of 'construction rules' in Prüst (1992) and Prüst et al. (1994), and as 'topic-based updating' in Asher (1993).

(2.5)

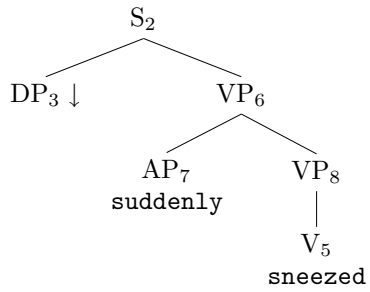

TAGs are not just employed for sentence analysis but also for discourse analysis. Both in (L)TAG for Discourse(Webber 1991; Schilder 1997; Gardent 1997a; Webber and Joshi 1998; Gardent and Webber 1998; Webber et al. 1999; Webber et al. 1999; Webber 2004)[5] and the Linguistic Discourse Model (Scha and Polanyi 1988; Prüst 1992; Prüst et al. 1994; Polanyi and van den Berg 1996; Polanyi et al. 2004) the adjunction operation plays a central role in modelling the incrementation process in the analysis of ongoing discourse, and in accounting for discourse syntactic constraints on interpretation.

In some variants of TAG, notably within Polanyi's Linguistic Discourse Model, another derivation operation is used beside adjunction, namely 'attachment' or 'extension'. *Extension* comes down to adding a subtree as a new daughter to a node in a tree which already has at least two daughters. The picture in (2.6) below illustrates this. In the Linguistic Discourse Model, extension is used when the conjugated discourse units are connected by a single relation of semantic parallelism, and thus, in the words of the authors,[6] are semantically strongly connected.

(2.6)

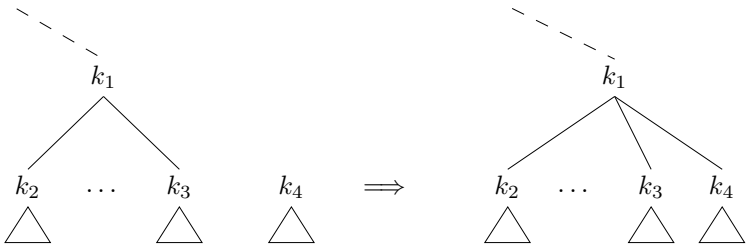

Much of the work on TAG for natural language concerns so-called Lexicalised TAG, see for example Schabes (1990), Joshi and Schabes (1991), and Abeillé (1993), Abeillé (1994). A *lexicalised TAG*, or *LTAG*, is a TAG in which each elementary tree structure is systematically associated with an overt lexical item called 'the anchor'. More precisely, a *structure*, elementary or derived, is lexicalised if there is at least one overt lexical item that appears in it. If more than

[5]Webber (2004) is a useful survey paper, setting pointers to other papers and approaches, and clarifying Webber's point of view on the sentence-discourse interface.

[6]Prüst et al. (1994), pp. 299, 300.

one lexical item appears, either one lexical item is designated as the anchor, or a subset of the lexical items local to the structure are designated as a 'multi-component anchor'. A *grammar* is lexicalised if it consists of a finite set of lexicalised structures and an operation or operations for composing the structures.

Usually, the *anchor* of an elementary tree instantiates its syntactic head. It determines what nodes and syntactic labels are present in the elementary tree in which it figures; these nodes can be said to be *anchored* to it. An elementary tree represents the *domain of locality* of its anchor, a domain over which syntactic and semantic constraints can naturally be stated. Because of the potential use of adjunction, this domain may be stretched or extended. Thus, LTAGs are characterised by an 'extended domain of locality' as compared to context-free phrase structure grammars.

2.3 Description Grammars

Although TAG is more powerful than context-free grammar, it has been argued that it still cannot handle some natural language phenomena very well. This is claimed, for example, with respect to certain aspects of scrambling in German (Rambow 1994a; Rambow 1994b), aspects of extraposition (Weir 1988; Kroch and Joshi 1987), and a uniform treatment of complements (Kroch 1987; Abeillé 1988a; Abeillé 1988b; Rambow et al. 1995a). In response to the last problem a variant of TAG called 'D-Tree Grammar' is developed in Rambow et al. (1995a), building on Vijay-Shanker (1992).

Both in Rambow's 'D-Tree Grammar' and Vijay-Shanker's rephrasing of TAG in terms of 'quasi-trees' (see below), tree descriptions or constraints on trees are used instead of trees, and the relation of immediate dominance connecting the nodes of a tree is relaxed. We shall call grammar formalisms which in one way or another share these two features *description grammars*. Beside D-Tree Grammar and TAG with quasi-trees, the term covers, among more, Logical Description Grammars (Muskens 1995; Muskens and Krahmer 1998; Muskens 2001), Radical Underspecification (Pinkal 1996), Constraint Language for Lambda-Structures (Egg et al. 1998; Egg et al. 2001), Tree Description Grammars (Kallmeyer 1999), and description grammar for discourse as proposed in Gardent and Webber (1998), Duchier and Gardent (1999), and Van Leusen and Muskens (2003).

The central characteristic of these grammar formalisms is that they generate descriptions of linguistic tree structures rather than the structures themselves. Either implicitly or explicitly there are two levels of analysis: the level of constraints or descriptions, and the level of the described linguistic structures. Since descriptions can be partial, they can serve to underspecify aspects of the structures they describe. In principle, any feature of the linguistic structure can be underspecified. Like TAGs, description grammars offer an extended domain of locality, but they have a greater generative capacity than TAGs.

The idea to use *descriptions* of trees in natural language processing can be traced back to Marcus et al. (1983). Marcus et al. are concerned with deterministic natural language parsing. They adopt the hypothesis that natural language can be analysed by a parsing process which builds a syntactic analysis *indelibly*, that is, any structure built by the parser is part of the correct analysis of the input. They observe, however, that when the indelible syntactic analysis takes the form of a tree, there are certain linguistic constructs which terminally confuse the parser, although they are interpretable to a human agent.

In order to solve this problem, they propose to use 'D-Theory', as they call it. Instead of trees, the parser should manipulate descriptions of syntactic structures or trees. The two characterising features of 'D-theory' are that it specifies 'dominance' rather than 'immediate dominance' between nodes, and that it uses names to refer to nodes. These node names function like variables: they can refer to any node that fits the description in as far as it is specified. They do not name a particular entity, as personal names do. *Immediate dominance* is a primitive relation between nodes (irreflexive, intransitive, asymmetric). A node i *dominates* a node j iff $i = j$ or there is a node k such that i dominates k and k immediately dominates j. By stating only dominance in a description, relations of immediate dominance existing in the tree described can be underspecified. This allows for deterministic parsing where otherwise some form of backtracing would be required.

Marcus et al. formulate a cognitive requirement for their approach: for D-theory to be psychologically valid, it must be the case that just those constructions which garden-path a D-theory parser garden-path people as well. This requirement is later referred to as the *deterministic hypothesis* in Duchier and Gardent (1999): cases of attachment ambiguity that require no conscious reanalysis from the language user must be processed deterministically.

The idea to use description theory for parsing is taken up in Vijay-Shanker (1992). In that paper a reinterpretation of TAGs as grammars which manipulate descriptions of trees, called 'quasi-trees', is proposed. Quasi-trees feature pairs of quasi-nodes which carry the same syntactic label and which are related through dominance: 'top' dominates 'bottom'. In building up a quasi-tree for a sentence, a pair of quasi-nodes may be identified, or it may be 'pushed apart' by intervening nodes.

Quasi-trees can be visualised by means of pictures that are similar to trees, but which should not be confused with them. In these pictures, immediate dominance is indicated by uninterrupted lines and dominance by dashed lines. For example, the quasi-trees describing the elementary trees from which the sentence 'John suddenly sneezed' is built up can be pictured as follows.

(2.7)

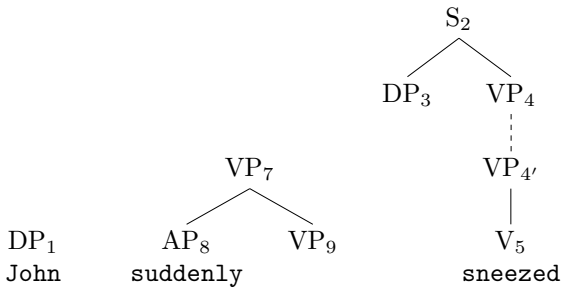

Elementary quasi-trees are put together through *subsertion* and *sister-adjunction*. Subsertion generalises over the substitution and the adjunction operation in TAG. Sister-adjunction has the same effect as extension in TAG.[7] Both these operations basically consist of the addition of dominance links between nodes. In our example, subserting the quasi-trees for 'John' and 'suddenly' into the quasi-tree for 'sneezed' results in the following quasi-tree.

(2.8)

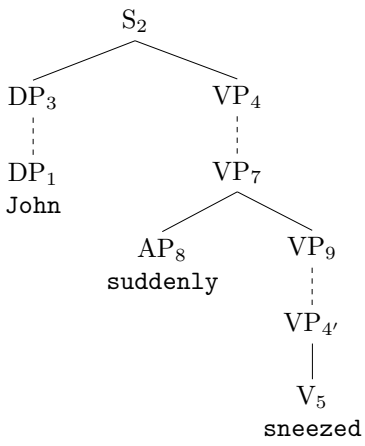

The set of *structures derived* by the quasi-tree grammar can be obtained by taking the 'circumscriptive reading' of the domination relation indicated in the quasi-trees composed with subsertion and sister-adjunction. This is done by identifying top and bottom nodes that are linked by dominance edges. For example, the circumscriptive reading of the quasi-tree in (2.8) is obtained by identifying node 1 and node 3, node 4 and node 7, and node 9 and node $4'$. Notice that a dominance relation may stretch over an arbitrary number of nodes linked by immediate dominance. In principle a quasi-tree describes an infinite number of tree structures. The tree structure obtained by taking the 'circumscriptive reading' of the domination relation in a quasi-tree is the *minimal* tree compatible with it.

[7] Actually these two operations do not feature yet in Vijay-Shanker (1992). We are grouping together here some follow up of Vijay-Shanker (1992), such as Rogers and Vijay-Shanker (1992), Rogers and Vijay-Shanker (1994), Rogers (1994), and Vijay-Shanker and Weir (1994), and closely related work on D-Tree Grammar, such as Rambow et al. (1995a) and Rambow et al. (1995b).

Vijay-Shanker's motivation for using descriptions is that it supports a straightforward embedding of (feature based) TAG in a unification framework. With standard TAG such an embedding is hindered by the fact that derivation is non-monotonic: the adjunction operation does not preserve the structural relations holding in the trees that are combined. In the quasi-trees approach, however, the derivation process is monotonic. Subsertion and sister-adjunction only add dominance links to elementary descriptions, and the circumscription operation only strengthens statements about dominance to identity statements.

This is not only computationally attractive but also linguistically interesting. Through the use of partial descriptions the syntax and semantics of the linguistic tree structures can be underspecified. Since natural language abounds in ambiguous phrases and constructs which underspecify the syntax and semantics of sentences, such as gapping and VP-ellipsis, it is desirable to have a grammar in which this can be expressed directly—a grammar which generates representations that underspecify the syntax or semantics of a sentence whenever it is in fact underspecified or ambiguous.[8] A grammar which produces partial descriptions of trees naturally does this.

The work by Marcus, Hindle and Fleck, and by Vijay-Shanker has been taken up by other authors, who propose their own variant of description grammar. They all use descriptions as underspecified representations of linguistic structures, but the grammars sometimes differ in their formal characteristics.

For example, the tree descriptions may be (sets of) quantified first order formulae as in Muskens and Krahmer (1998) and Muskens (2001) or nonquantified ones as in Kallmeyer (1999), Egg et al. (1998) and Duchier and Gardent (1999). The described tree structures may provide both syntactic and semantic information as in Muskens and Krahmer (1998), or they may be just semantic representations as in Egg et al. (1998). Description grammar is combined with compositional DRT as a semantic representation language in Van Leusen and Muskens (2003), and with (U)DRT in Cimiano and Reyle (2005). Blackburn and Bos (1999) and Bos (2004) combine λ-DRT with Hole Semantics, an underspecification formalism closely related to UDRT.

Furthermore, in most papers the described structures are regular tree structures, but in Blackburn and Gardent (1998) and Duchier and Gardent (1999) they are proposed to be 're-entrant' trees (Directed Acyclic Graphs). Kallmeyer defines 'synchronous local TDGs' and employs them in the syntax-semantics interface of her grammar (1999). The manner in which the minimal trees or 'solutions' or 'models' of a description are computed may vary: in Muskens and Krahmer (1998) and Muskens (2001) built-in properties of the grammar and the description of a sentence ensure that only minimal trees fit the description, in other approaches, for instance Duchier and Gardent (1999), sentence descriptions allow

[8]Such representations are commonly, perhaps confusingly, called 'underspecified representations'. The representations themselves are not underspecified, of course.

nonminimal verifying trees but only the minimal ones are computed.

There is also some variation with respect to the issues addressed, and the perspective from which they are viewed. The emphasis may be on computational, linguistic, or cognitive issues. In Marcus et al. (1983), tree descriptions are proposed to be used for computational purposes, notably, to allow for deterministic parsing of natural language. In Vijay-Shanker (1992), a description grammar, manipulating quasi-trees, is developed in order to define derivation as a monotonic process. In Erk et al. (2002) the computational properties of processing CLLS are investigated. Duchier and Gardent (1999) zooms in on the efficient computation of minimal models for description grammar allowing re-entrant trees. Blackburn and Bos (1999, 2005) and Bos (2004) provide an extended computational implementation and integrate automated inference techniques.

Taking up the deterministic hypothesis of Marcus et al. (1983), Gardent and colleagues address cognitive issues in the treatment of garden-path sentences and incremental discourse processing (Gardent and Webber 1998; Duchier and Gardent 1999). Interestingly, they argue that for a proper treatment of garden-path sentences more is needed than just description theory: underspecification must be combined with a preference system that highlights a specific reading corresponding to the hearer's currently preferred interpretation. Without this additional bias, description grammar would be too tolerant, and would not predict that in these cases some form of backtracking or revision of previous analysis takes place.

Finally, different authors concentrate on different linguistic constructs. Among others, Rambow et al. (1995a) discuss the treatment of complement clauses, and *wh*-movement. The treatment of quantifier-scope ambiguities, ellipsis and anaphora at sentence level is in focus in what might be called the Saarbrücken branch of description grammar, that is Pinkal (1996), Egg et al. (2001), and others. Scope ambiguities and lexical ambiguity are also addressed in Muskens and Krahmer (1998), Muskens (2001), and Cimiano and Reyle (2005). VP-ellipsis is in focus in Thompson (2002), who combines description theory with higher order unification. The constraining effect of ontological information on syntactic and semantic ambiguity is addressed in Cimiano and Reyle (2005). Various forms of structural and semantic ambiguity occurring at discourse level are treated in Gardent and Webber (1998) and Duchier and Gardent (1999). Garden-path sentences are discussed in both Marcus et al. (1983) and Gardent and Webber (1998). In Van Leusen and Muskens (2003) the interaction of anaphora and presupposition resolution with general background information at discourse level is addressed. Egg and Redeker (2006) focus on discourse structure per se.

2.4 LDG in Detail

The particular version of description grammar that will be employed in our framework of discourse interpretation is the one proposed in Muskens and Krahmer

(1998) and Muskens (2001). Following Muskens, we call the grammar a *Logical Description Grammar* or *LDG*. The name highlights the fact that the descriptions generated by the grammar are statements or sets of statements in a logical language, in this case classical type logic. The grammar is purely declarative. It models parsing as a deduction process; a parse results from the language user's reasoning about linguistic descriptions. The formalism is nonincremental, that is, all relevant linguistic information for the analysis of a sentence must be given at the outset of the reasoning process. In the next chapter, where we apply the formalism to the description of discourse structure, an incremental version of the grammar will be proposed.

In the following sections the LDG formalism will be discussed in detail. Section 2.4.1 explains the overall working of the grammar. In section 2.4.2, we go through an example which illustrates the parsing process and which shows how structural ambiguity is dealt with in an LDG. This introduces some further aspects of Muskens' approach in an informal way. Subsequently, we zoom in on the various components of the grammar formalism. Section 2.4.3 introduces the axioms or general descriptions of the grammar, section 2.4.4 is about the positive and negative polarities of nodes, in section 2.4.5 input descriptions are discussed, and section 2.4.6 is about lexical descriptions.

2.4.1 Overall Working of the Grammar

An *LDG system* $G = \langle \mathcal{G}, I \rangle$ consists of the following components.

- A *Logical Description Grammar* $\mathcal{G} = \langle \mathcal{A}, \mathcal{W} \rangle$, where
  - $\mathcal{A}$ is a set of axioms or general descriptions and
  - $\mathcal{W}$ a set of lexical descriptions (classifying and elementary)
- I, an infinite set of *input descriptions* δ, generated per processed utterance

Lexical descriptions describe the syntax and semantics associated with the lexemes of the language. They are comparable to the elementary trees in a TAG. Lexical descriptions come in two kinds: *classifying descriptions* and *elementary tree descriptions*. The first link a lexeme to its word class and a semantic value, the second associate a tree structure with a basic word class. Differing slightly from Muskens (2001) we will use the term lexical description only to indicate the combination of a classifying description and the appropriate elementary tree description. The *axioms* or *general descriptions* determine the structures that are described by the grammar, namely tree structures.

A full-fledged LDG can be expected to contain additional descriptions representing aspects of the specific *linguistic theory* incorporated in the grammar. As

opposed to the axioms, these descriptions would cover the 'variable', or theory-dependent part of the grammar formalism. The linguistic theory might be presented as a separate component of the grammar, called $\mathcal{T}$. As we will see, however, constraints of the linguistic theory can very often be incorporated in the lexicon $\mathcal{W}$, or else formulated in a few additional general descriptions, so that there is no necessity to adopt an new component $\mathcal{T}$.

In what we shall often call the *grammar system* or *grammar formalism* the grammar $\mathcal{G}$ is combined with an infinite set of *input descriptions* I, which are generated per sentence. Input descriptions sum up the lexical items in a sentence and specify the order in which they occur. Beside the grammar and the input descriptions an LDG system may be assumed to contain so-called *integrity constraints*: constraints of a logical or linguistic nature on the form of input descriptions and lexical descriptions. Being meta-level constraints, we have not represented them in the system.

All descriptions in the grammar system, whether input, lexical, or general descriptions, are stated as quantified formulae in first order predicate logic. The grammar is a set of logical statements, a logical theory representing the linguistic knowledge of the language user. This facilitates reasoning with descriptions. In particular, parsing can be modeled as logical deduction given $\mathcal{G}$. A description grammar $\mathcal{G}$ is *adequate* for a given language if, for any grammatical sentence u in the language and for no ungrammatical one, the input description of the sentence δ_u is consistent with $\mathcal{G}$. *Sentence interpretation* is modeled as follows in LDG.

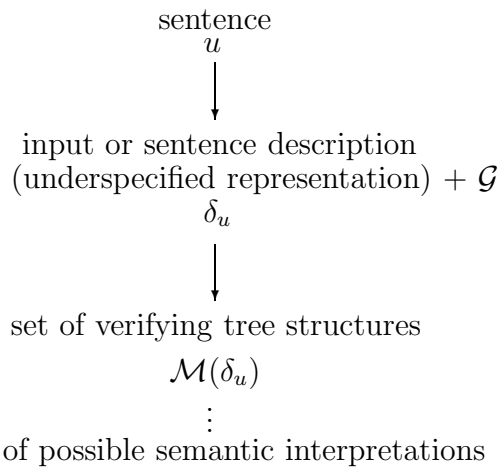

Figure 2.1: Sentence interpretation in LDG

First, the language user constructs an input description of the sentence and selects the relevant lexical descriptions from $\mathcal{W}$. Though the term is not used in Muskens (2001), we will refer to the whole of this as the *sentence description*. It represents the language user's linguistic knowledge of the sentence and the words occurring in it. Sentence descriptions can be partial, and so, syntactic or semantic features

of the sentence may be underspecified. Sentence descriptions instantiate what might be called the level of underspecified representation in the grammar system.

Next, the language user may reason about the tree structures that 'fit' the sentence description given $\mathcal{G}$, thereby *parsing* the sentence. Parsing is a deduction process in which the minimal trees or models that verify a sentence description are determined.[9] The resulting set of trees instantiates a second level of analysis in the grammar, the level of fully specified linguistic structures.

Reasoning about the alternative sets of node identifications which the sentence description allows for given the axioms in $\mathcal{G}$, the language user may arrive at a set of alternative, more specific descriptions. Each of these alternatives is validated by a particular tree structure.[10] The respective tree structures are precisely those that validate the original description. Each of them corresponds to a particular disambiguation of the sentence. When the description of a sentence is not validated by any tree structure, the sentence does not belong to the language generated by the grammar.

Finally, to obtain the semantic interpretation of the sentence, the language user may collect the semantic values at the root nodes of the tree structures validating the sentence description. The interpretation of these values produces the possible interpretations or readings of the sentence. In fact, semantic interpretation is an integrated part of the parsing process: reasoning about the sentence description involves both reasoning about the tree structures that satisfy it and reasoning about the potential meaning of the sentence, as represented at the roots of the tree structures. Like syntactic features, semantic values can be underspecified.

In the following section the parsing process is illustrated with an example. After that, we zoom in on the different components of the grammar and the parsing-as-deduction method. Section 2.4.3 discusses the axioms of the grammar, section 2.4.4 is about the pairing of 'positively' and 'negatively marked' nodes in the parsing process—not discussed so far—, in section 2.4.5 input descriptions are examined, and section 2.4.6 is about lexical descriptions.

2.4.2 Parsing an Ambiguous Sentence: An Example

Consider the sentence

(2.9) *John saw a man with a telescope.*

[9]The tree structures are indeed the models which validate the descriptions according to classical model-theory. In order to avoid confusion between the models of descriptions, and the models commonly featuring in semantic interpretation (also represented in the schema in 2.1) we prefer to use the terms 'trees' or 'tree structures'.

[10]To be precise, each alternative description is validated by the tree structures in a single equivalence class of structures.

This is a classic case of syntactic ambiguity: the modifying PP 'with a telescope' may be attached to the NP, as in the syntactic tree in (2.10) below, or to the VP, as in tree structure (2.11). This results in, respectively, the reading of the sentence in which there was a man with a telescope whom John saw, or the reading in which John saw a man by means of a telescope.

(2.10)

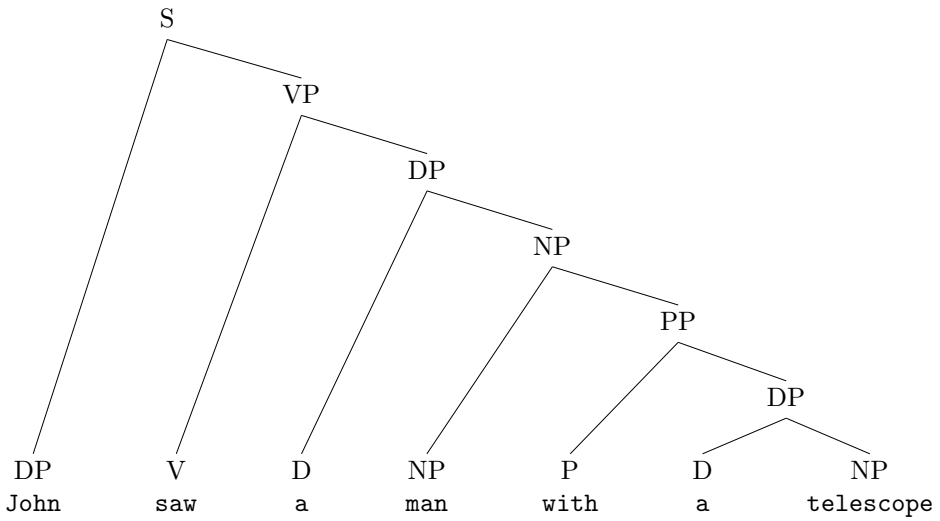

(2.11)

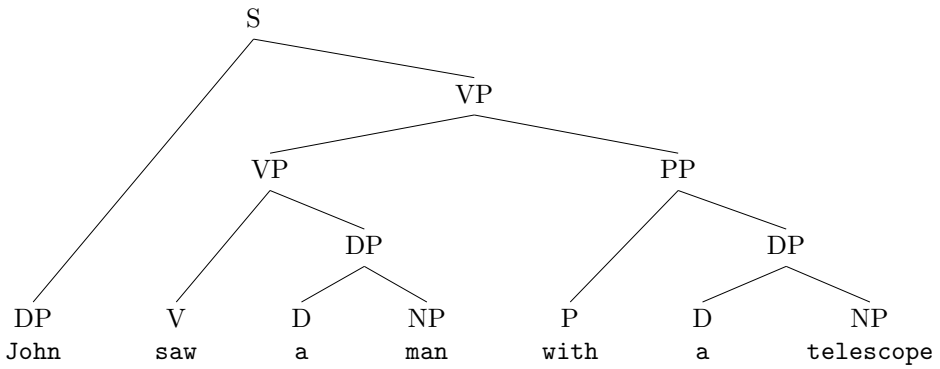

In a standard TAG, it is not possible to generate a single syntactic representation which underspecifies the ambiguity, thus reflecting that the sentence may be interpreted either way. This is different with description grammars. In LDG, the two trees in (2.10) and (2.11) feature only at the second level of analysis in figure 2.1, the level of fully specified linguistic structures. Processing the sentence results in a possibly partial logical description of the sentence, the first level of analysis in figure 2.1. The ambiguity of the sentence is captured in the sentence description through underspecification of the syntactic structure.

As in section 2.3, the sentence description of (2.9) can be represented in the form of a picture; see (2.12) below. As before, uninterrupted lines represent the relation of immediate dominance and dashed lines represent dominance over an arbitrary, underspecified length. The left-right ordering between sister nodes and between lexical nodes in the picture reflects the precedence relation between these

nodes. Furthermore, the picture now figures a root node r and a variable category label XP, which can match either NP or VP.[11]

(2.12)

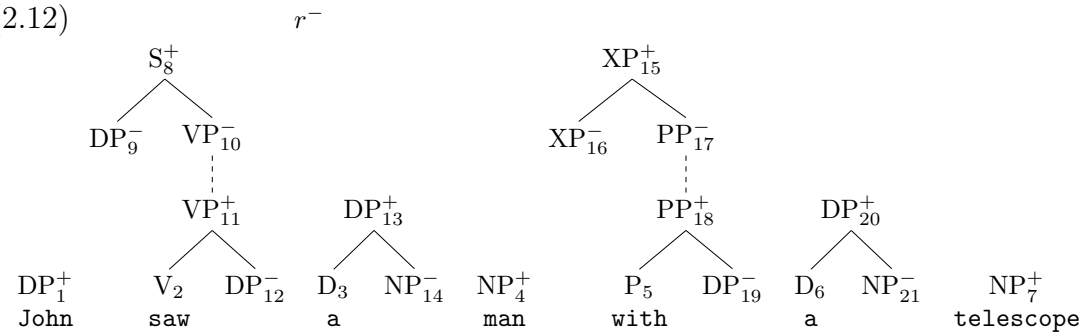

The picture in (2.12) shows an important new feature, namely the decoration of certain nodes with $^+$ and $^-$ polarities. The trees or models that verify a sentence description can be obtained by pairing off positively and negatively marked nodes in a one-to-one fashion and identifying the nodes thus paired. The axioms and integrity constraints of the grammar ensure that this is so, see section 2.4.4 below. They also enforce that the resulting structure is a linguistic tree, that the order of the lexical elements in the structure reflects their order in the sentence, and that category information linked to the nodes is respected that is nodes can only be identified when they carry the same syntactic label.

Given the sentence description, the sentence can be parsed. Remember that parsing is a deduction process. Using his grammatical knowledge, the language user may deduce an output description which is more explicit about the models or tree structures that satisfy it. In particular, he will reason about the mapping of positively and negatively marked nodes and the resulting node identifications.

In the case of the description in (2.12) he might reason as follows. Consider the negatively marked node named 21. It can be seen that $21 = 7$ can be derived: 21 must match with a positively marked node. All nodes named other than 4, 15 and 7, are excluded on the basis of clashing category information; but of these three 4 and 15 are out because they precede 6 while 21 is preceded by 6. Now only one candidate is left, namely 7, so 21 must be identified with 7. In much the same way $19 = 20$, $12 = 13$, $9 = 1$, $17 = 18$ and $8 = r$ is derived. This produces the following description as a possible intermediate result.

[11]The labeling 'XP' on nodes 15 and 16 in (2.12) is just a way of representing the fact that the syntactic label of these nodes is underspecified, though they must carry the same label. On the assumption that a modifier headed by the preposition 'with' may attach to either an NP or a VP, XP may be either an NP or a VP. This cannot be expressed very well by means of a picture, but in the lexical description of the preposition it can be spelled out. It should contain the following statement, in which ℓ is a function which assigns syntactic labels to nodes: $(\ell(15) = vp \vee \ell(15) = np) \wedge \ell(16) = \ell(15)$

(2.13)

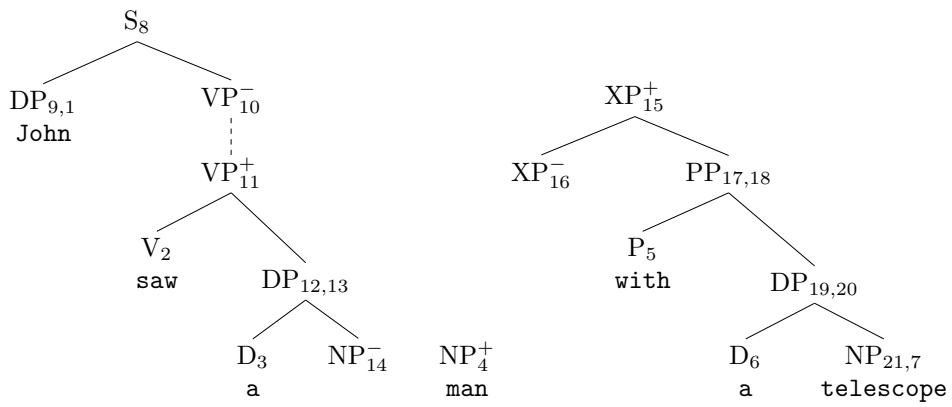

At this point there is a choice between the identification of node 16 with node 11 or with node 4. Now it cannot be deduced *which* identification must be made. This reflects the ambiguity in the sentence which is being parsed. However, when each of the two possibilities is pursued we find the two tree structures which validate the description, and which represent the two possible interpretations of the sentence.

If node 16 is identified with node 11 then given the general descriptions it logically follows that $10 = 15$, and $14 = 4$. This produces one of the possible models of the description, namely the tree-structure introduced earlier in (2.11). Alternatively, if node 16 is identified with node 4, it can be deduced that $14 = 15$, and $10 = 11$. This results in the other possible model, namely the tree structure in (2.10). Through the combined force of the lexical, input and general descriptions these two are ensured to be the only models of (2.13). The sentence description captures the ambiguity of the sentence in a single representation through the underspecification of dominance relations.

2.4.3 General Descriptions

As stated in section 2.4.1 an LDG consists of a set of axioms or general descriptions, a set of lexical descriptions, and input descriptions, generated per sentence. The general descriptions of the grammar characterise the basic properties of the structures described. They ensure that the structures are always tree-like. Given the additional constraining force of the lexical descriptions and the input descriptions the structures described by the grammar must be linguistic trees.

The set of general descriptions proposed in (Muskens 2001) comprises fourteen axioms, the first ten of which concern dominance and precedence in tree structures. In the description language of the grammar nodes are indicated by $k, k_1, k_2, \ldots$. Furthermore $\prec$ denotes *linear precedence*, $\lhd$ denotes *immediate dominance*, $\lhd^*$ denotes *dominance*, and $\lhd^+$ *proper dominance*. A node k_1 dominates a node k_2 if k_1 properly dominates k_2, or if k_1 is identical to k_2.

$\mathcal{A}1$ $\forall k\, r \lhd^* k$

$\mathcal{A}2$ $\forall k \neg k \lhd^+ k$

$\mathcal{A}3$ $\forall k_1 k_2 k_3 [[k_1 \lhd^+ k_2 \wedge k_2 \lhd^+ k_3] \rightarrow k_1 \lhd^+ k_3]$

$\mathcal{A}4$ $\forall k \neg k \prec k$

$\mathcal{A}5$ $\forall k_1 k_2 k_3 [[k_1 \prec k_2 \wedge k_2 \prec k_3] \rightarrow k_1 \prec k_3]$

$\mathcal{A}6$ $\forall k_1 k_2 [k_1 \prec k_2 \vee k_2 \prec k_1 \vee k_1 \lhd^+ k_2 \vee k_2 \lhd^+ k_1 \vee k_1 = k_2]$

$\mathcal{A}7$ $\forall k_1 k_2 k_3 [[k_1 \lhd^+ k_2 \wedge k_1 \prec k_3] \rightarrow k_2 \prec k_3]$

$\mathcal{A}8$ $\forall k_1 k_2 k_3 [[k_1 \lhd^+ k_2 \wedge k_3 \prec k_1] \rightarrow k_3 \prec k_2]$

$\mathcal{A}9$ $\forall k_1 k_2 [k_1 \lhd k_2 \rightarrow k_1 \lhd^+ k_2]$

$\mathcal{A}10$ $\forall k_1 k_2 k_3 \neg [k_1 \lhd k_3 \wedge k_1 \lhd^+ k_2 \wedge k_2 \lhd^+ k_3]$

Axiom $\mathcal{A}1$ states that the root r dominates all other nodes. $\mathcal{A}2$–$\mathcal{A}5$ ensure that $\lhd^+$ and $\prec$ are strict partial orderings. The *Exhaustivity* property, $\mathcal{A}6$, says that any two nodes stand either in a precedence relation or in a relation of dominance.[12] The *Inheritance* property, laid down in $\mathcal{A}7$ and $\mathcal{A}8$, says that precedence relations are inherited from ancestor nodes. Finally, $\mathcal{A}9$ and $\mathcal{A}10$ constrain the relation between immediate dominance ($\lhd$) and proper dominance ($\lhd^+$).

The description language further contains the *labeling function* ℓ. The labeling function assigns syntactic category labels to nodes, as for example in $\ell(n) = vp$. To ensure that nodes are labeled uniquely, it must be prevented that category label names corefer. This is what the following axiom does.

$\mathcal{A}11$ $c_1 \neq c_2$, if c_1 and c_2 are distinct label names

Three further axioms will be discussed in the next section.

2.4.4 Positive and Negative Anchoring

It is a specific property of the description grammar formalism as it is formulated in Muskens (2001) that nodes in descriptions are either positively or negatively marked, or not marked at all, in which case they are called *saturated*. Saturated nodes are of both $^+$ and $^-$ polarity; they are unmarked in order to keep the pictures of descriptions transparent. To illustrate, in the following description of the sentence 'John sneezed', the nodes 1, 3, and 6 are positively marked, the nodes 4 and 5 and the root are negatively marked, and 2 and 7 are saturated.[13]

[12] Furthermore, it follows from $\mathcal{A}1$–$\mathcal{A}8$ that every pair of distinct nodes that is not related through dominance is related through precedence. This is known as the Exclusivity property, see for instance Partee et al. (1990), ch. 16 p. 442.

[13] Negatively marked nodes correspond to what are sometimes called 'open nodes' in other description grammar formalisms, compare for example Duchier and Gardent (1999). Positively marked, and saturated nodes correspond to 'closed nodes'.

(2.14) r^-

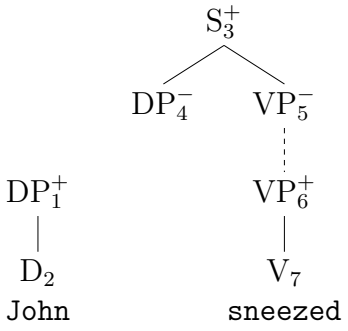

The positive and negative marking of nodes plays a central role in obtaining the tree structures that satisfy a given description. The verifying trees are obtained by mapping positively and negatively marked nodes one-to-one. Negatively marked nodes may attach to material below them, positively marked nodes may attach to material above them, and saturated nodes are not required to be identified with any other node.

All of this is regulated by the axioms $\mathcal{A}12$–$\mathcal{A}14$, which will be listed below. Rather than talk about the positive or negative *marking* of nodes in descriptions of trees, however, these axioms talk about the positive or negative *anchoring* of nodes to their lexical anchor. The notion 'lexical anchor' already popped up in our discussion of the lexicalisation constraint in TAG in section 2.2. The constraint requires that each elementary tree in a TAG is associated with (at least) one lexical item, called its *anchor*. The anchor determines what syntactic constituents feature in the elementary tree. All nodes in the elementary tree are anchored to the designated lexical item or to the node at which it sits.

In LDG the lexicalisation constraint is implemented as follows. For every lexical description in $\mathcal{W}$ it holds that all the nodes in the description are anchored to the lexical head of the described tree. If a constituent in the described elementary tree is, in the sense of X-bar theory, a projection of the lexical head, the corresponding node is positively anchored to it. The intuitive idea behind this is that a positively anchored node in an elementary tree is introduced, or contributed by the anchor. On the other hand, if a constituent is subcategorised for by the lexical head, the corresponding node is negatively anchored to it. A negatively anchored node in an elementary tree represents a constituent which is required or 'consumed' by the lexical head. It must be produced by other elementary trees.

To clarify the above, consider the following representation of the sentence description in (2.14), in which the anchoring relations expressed by the positive and negative marking in (2.14) are visualised by arrows, and the anchors are indicated by $\diamond$.

(2.15)

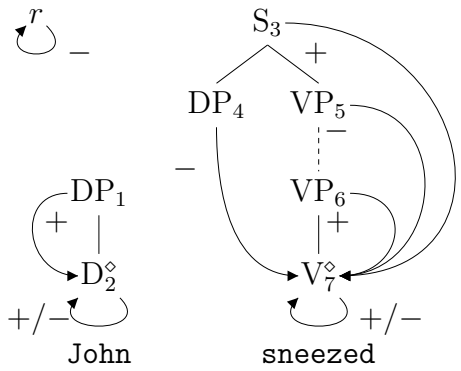

Almost all of the anchoring relations in this description are introduced by the lexical descriptions of *John* and *sneezed*, the only exception is the negative anchoring of the root. The anchoring relations play an essential role in deducing node identifications in a given sentence description. The deduction process is constrained by the following three general descriptions, which we shall call the *anchoring* axioms of the grammar. The predicate *lex* says that a node is lexical, and the functions α^+ and α^- produce, respectively, the positive and negative anchor of a node. For example $\alpha^+(k_1) = k_2$ says that k_1 is positively anchored to k_2.

$\mathcal{A}12$ $\forall k\, lex(\alpha^+(k))$

$\mathcal{A}13$ $\forall k\, [k = r \vee lex(\alpha^-(k))]$

$\mathcal{A}14$ $\alpha^-(r) = r$

Axiom $\mathcal{A}12$ requires that the positive anchor of a node is lexical and $\mathcal{A}13$ requires the same for negative anchoring, except for the root. The root is negatively anchored to itself ($\mathcal{A}14$). Intuitively, this means that the root creates the open slot it fills itself: there is no element which subcategorises for the root.

Together, the axioms $\mathcal{A}12$–$\mathcal{A}14$ require that each node in a tree structure is both positively anchored to a lexical node, and negatively anchored to a lexical node or the root. The positive and the negative anchor can but need not be identical. If the anchoring axioms are to be satisfied, nodes which in a given lexical or derived description are only specified to be positively anchored must be identifical to nodes that are only specified to be negatively anchored, and vice versa. Thus positively marked nodes are mapped one-to-one to negatively marked ones. The anchoring axioms help to ensure that the trees verifying a sentence description are the *minimal* ones verifying it. The distinction between positive and negative anchoring relations is essential for this.

In the case of example (2.15), the anchoring axioms enforce the identification of node 1 and 4, node 5 and 6, and node 3 and the root. Note that in the resulting tree, see (2.16) below, node 1 (or 4) is positively anchored to node 2 but negatively anchored to node 7. This reflects that the DP is contributed or produced by the proper name *John* and at the same time functions as an argument of the predicate headed by *sneezed*.

(2.16)

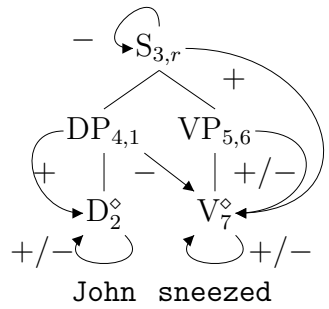

To end this discussion of positive and negative anchoring, notice that the anchoring axioms cannot be satisfied relative to individual lexical descriptions. The axioms regulate the manner in which lexical description may combine to form a sentence description, hence they can only be sensibly applied at sentence level. Even then, it must be so that the input description of a sentence triggers 'enough' lexical tree descriptions. If not, either $\mathcal{A}12$ or $\mathcal{A}13$ is violated and inconsistency can be deduced. It would follow that the sentence is ungrammatical.

2.4.5 Input Descriptions

The description grammar outlined above is a sentence grammar. For any given sentence, it produces an input description which covers the complete sentence. Input descriptions sum up the lexical nodes appearing in the sentence, and specify the linear order of these nodes and the lexical items attached to them. Furthermore, input descriptions contain what we shall call a *circumscription clause*. The clause says that there are no other lexical nodes than the ones named in the input description. To illustrate, in (2.18) we give the input description of the sentence (2.17). The circumscription clause is underlined.

(2.17) *John suddenly sneezed.*

(2.18) $\exists k_1 k_2 k_3 (k_1 \prec k_2 \prec k_3$
$\wedge \textsc{john}(k_1) \wedge \textsc{suddenly}(k_2) \wedge \textsc{sneezed}(k_3)$
$\wedge \underline{\forall k(lex(k) \leftrightarrow (k = k_1 \vee k = k_2 \vee k = k_3))})$

It is one of the integrity constraints of the grammar framework that the order of the lexical nodes in an input description reflects the order of the words in the sentence it represents. Presumably, another integrity constraint—though not one stated in Muskens (2001)—is that there is a one-to-one correspondence between the words in a sentence and the lexical nodes occurring in its input description. This ensures that the input description of a sentence triggers 'enough' lexical tree descriptions, so that, in principle, the anchoring axioms can be satisfied. Finally, the presence of a circumscription clause in the input description is another integrity constraint.

The circumscription clause plays a central role in ensuring that the described trees are finite, and more specifically, that they are the *minimal* ones that may satisfy the description. The clause requires the lexical nodes in any tree that may satisfy the input description to be one of a particular finite set, namely those that are named in the input description. Similarly, the lexical descriptions triggered by the lexical nodes in the input description specify and circumscribe the nonterminal nodes that are positively or negatively anchored to each of the lexical nodes, and this is a finite set too. In this manner, the maximal number of nodes (both terminal and nonterminal) in a tree is determined: there can be no more nodes in a verifying tree than the ones that are explicitly named in the input description and the relevant lexical descriptions.

The LDG formalism presentend here is purely declarative and nonincremental: input descriptions are not built up word by word but generated at once for the whole sentence. Logical deduction is applied to complete input descriptions. The circumscription clause, formulated as it is, would cause problems in an incremental set-up of the grammar. This can easily be seen by considering what would happen if the grammar formalism would be used to process discourse incrementally. Suppose a discourse is a sequence of sentences. We would then process the sentences one after another, with each incrementation step adding a new input description to the ones already generated. Each individual input description in the resulting set or discourse description would contain its own circumscription clause, saying that there are no other lexical nodes than the ones described in it. This would cause the discourse description to be inconsistent. Thus, in an incremental version of the formalism, circumscription must be implemented in a different way. We will return to this issue in the next chapter, in which an incremental variant of LDG will be proposed.

2.4.6 Lexical Descriptions

So far we have been talking about the lexicon of the description grammar as if it consists of a single type of descriptions, namely lexical descriptions. In fact, the lexicon, as it is defined in (Muskens 2001), contains two kinds of descriptions, namely classifying descriptions and elementary tree descriptions.

Classifying descriptions specify, for every lexical item, its semantic value and the word class it belongs to. In the case of example (2.17) the following descriptions apply. They classify 'John' as a proper noun (pn), 'suddenly' as an adverb (adv), and 'sneezed' as an intransitive verb (iv). The function σ gives the semantic value of a node. More precisely, σ stands for a family of functions σ^α, which associate nodes with semantic values of type α. For simplicity, the superscripts are suppressed in our descriptions.

(2.19) a. $\forall k[\textsc{john}(k) \rightarrow (pn(k) \wedge \sigma(k) = john)]$

b. $\forall k[\textsc{suddenly}(k) \rightarrow (adv(k) \wedge \sigma(k) = suddenly)]$
c. $\forall k[\textsc{sneezed}(k) \rightarrow (iv(k) \wedge \sigma(k) = sneezed)]$

Elementary tree descriptions associate a tree structure, including syntactic and semantic labels, with each basic word class. An elementary tree description describes the syntax and semantics contributed by an arbitrary lexical item in the relevant word class. While classifying descriptions provide the basic semantic values of lexical nodes, elementary descriptions provide 'recipes' for the composition of the complex semantic values of nonterminal nodes. The semantic representation language used in Muskens (2001) is classical type logic.[14] In principle, however, the description grammar formalism does not enforce the use of a particular semantic representation language, and a different representation language could be chosen. In our LDG for discourse compositional DRT will be employed, see chapter Four.

The grammar might associate the following elementary tree desciptions with the word classes *pn*, *adv* and *iv*. In the descriptions $\triangle(k, k', k'')$ abbreviates the conjunction $k \lhd k' \wedge k \lhd k'' \wedge k' \prec k''$. Furthermore, $t \overset{+}{\hookleftarrow} \{t_1, \ldots, t_n\}$ and $t \overset{-}{\hookleftarrow} \{t_1, \ldots, t_n\}$ abbreviate the information that $t_1, \ldots, t_n$ are exactly the elements that are, respectively, positively and negatively anchored to t.

(2.20) $\forall k[pn(k) \rightarrow (\ell(k) = dp \wedge k \overset{+}{\hookleftarrow} \{k\} \wedge k \overset{-}{\hookleftarrow} \emptyset)]$

(2.21) $\forall k[adv(k) \rightarrow \exists k_2 k_3[\ell(k) = ap \wedge \ell(k_2) = vp \wedge \ell(k_3) = vp \wedge$
$\triangle(k_2, k, k_3) \wedge \sigma(k_2) = \sigma(k)(\sigma(k_3)) \wedge k \overset{+}{\hookleftarrow} \{k, k_2\} \wedge k \overset{-}{\hookleftarrow} \{k, k_3\}]]$

(2.22) $\forall k[iv(k) \rightarrow \exists k_2 k_3 k_4 k_5[\ell(k) = v \wedge \ell(k_2) = s \wedge \ell(k_3) = dp \wedge$
$\ell(k_4) = vp \wedge \ell(k_5) = vp \wedge \triangle(k_2, k_3, k_4) \wedge k_4 \lhd^* k_5 \wedge$
$k_5 \lhd k \wedge \sigma(k_5) = \sigma(k) \wedge \sigma(k_2) = \sigma(k_4)(\sigma(k_3)) \wedge$
$k \overset{+}{\hookleftarrow} \{k, k_2, k_5\} \wedge k \overset{-}{\hookleftarrow} \{k, k_3, k_4\}]]$

Notice that the elementary tree descriptions for proper names (2.20) and adverbs (2.21) have been simplified, for purely expository reasons. They do not contain a node which is labeled by the syntactic category label corresponding to the basic word class. The nodes labeled 'DP' and 'AP' act as lexical anchors.

We shall continue to use the term *lexical description* for the combination of the classifying description of a particular lexical item with the elementary tree description of the appropriate word class. To illustrate, for the lexical item

[14] It is type logic with an extension. Muskens embeds predicate logic (PL) in classical type logic in order to be able to model variable binding in quantificational contexts after the binding machinery used in the Tarski truth definition of PL. The resulting language contains formulae of type 'register' (comparable to discourse referents) and of type 'state'. A more extended explanation of the embedding of PL in classical type logic than the one provided in Muskens (2001) can be found in Muskens (1996).

'suddenly', the lexical description resulting from conjoining (2.19b) and (2.21) can be graphically represented as follows.

(2.23)

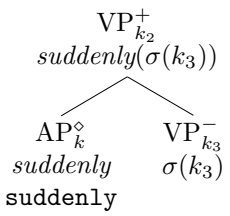

In this picture node names appear as subscripts. Semantic representations are represented in italics; note that $\sigma(k)\sigma(k_3)$, occurring in the elementary tree description for adverbs, has been instantiated to *suddenly*$(\sigma(k_3))$. Lexical descriptions for the other lexical items in 'John suddenly sneezed' can be constructed in the same way.

As was explained in the preceding sections, *parsing* is a deduction process in LDG. Parsing the sentence 'John suddenly sneezed' consists in collecting the input description of the sentence (2.18) and the classifying- and elementary tree descriptions describing the lexical elements in it, and reasoning about the tree structures that verify them given the axioms of the grammar. As a preliminary, witnesses are chosen for the nodes described in the input description. Intuitively, this implements the idea that the language user is parsing a particular occurrence of the sentence and the lexical items it contains. Whenever he comes across new nodes in the reasoning process fresh witnesses are taken. To illustrate, he may take witnesses 1, 2, 3, and 4 for the nodes in (2.18), which produces

(2.24) $1 \prec 2 \prec 3 \wedge \textsc{john}(1) \wedge \textsc{suddenly}(2) \wedge \textsc{sneezed}(3)$
$\wedge \forall k(lex(k) \leftrightarrow (k = 1 \vee k = 2 \vee k = 3))$

In the other descriptions, node names can subsequently be substituted by the appropriate witnesses. Taking input and lexical descriptions together this results in

(2.25)

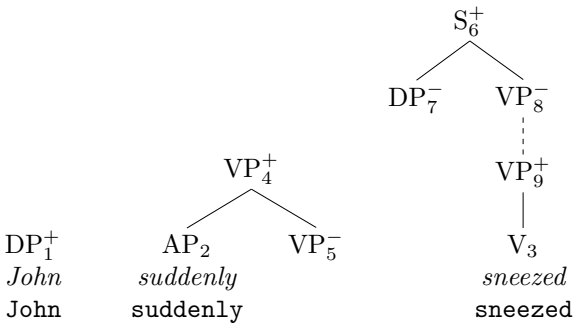

Using this description (or set of descriptions) as a premise to his grammar, the language user can now deduce the node name identifications $7=1$, $5=9$, and $8=4$.

See Muskens (2001) for a detailed discussion of the reasoning process involved, made relative to a similar example. The node name identifications result in a description which is verified by a single tree structure, namely

(2.26)

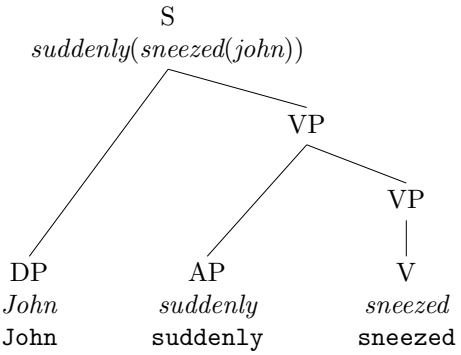

The axioms of the grammar do not allow for the deduction of any alternative set of node identifications. Thus, the sentence is predicted to be unambiguous.

In case a sentence is ambiguous (as is, for example, (2.9) that was discussed in section 2.4.2) there will be more than one model or tree structure that satisfies its description. Notice, however, that although the alternative models that satisfy a description in case of ambiguity can be determined by means of logical deduction, a partial description cannot be made *more specific* by logical deduction alone. Ambiguity and underspecification can only be resolved by adding more information, that is, by increasing the premises in the deduction process. This is in accordance with our intuitive understanding of the matter: an ambiguous statement can only be disambiguated by taking into consideration additional information or 'context'.

2.5 Conclusion

LDG combines grammar with reasoning and underspecification theory. The grammar is a logical theory, representing the language user's grammatical knowledge. Parsing, or more in general, linguistic interpretation, is a reasoning process: on perceiving a sentence the language user constructs a logical description of the words occurring in it and the order in which they occur. Given his grammatical knowledge, the language user can add descriptions of the lexemes in the sentence, and reason about the linguistic tree structures that fit or verify the sentence description. Different verifying tree structures represent different potential readings of the sentence.

Logical descriptions function as so-called underspecified representations. Since descriptions can be partial, syntactic and semantic features of sentences may be underspecified. A consequence of this is that various forms of underspecification and ambiguity occurring in natural language expressions can be modeled directly.

From a cognitive linguistic point of view there is a strong intuitive appeal in this: descriptions never contain more information than can be justified on linguistic grounds. On the basis of his grammatical knowledge the language user may reason about a description and find out more about the linguistic structures that satisfy it. He will only disambiguate an ambiguous sentence or construction when through background or world knowledge, or through subsequent conversation, he obtains more information about it, so that it can be deduced that there is only a single linguistic structure satisfying the description.

In a computational setting, it is an advantage that the derivation or incremental parsing of linguistic expressions can be modeled as a monotonic, nondestructive process. During the parsing process no information is retracted from the sentence description. An interesting cognitive claim relating to this feature is the deterministic hypothesis, which says that cases of attachment ambiguity that require no conscious reanalysis from the language user must be processed deterministically. As is argued in Marcus et al. (1983) and Gardent and Webber (1998), however, to account for sentences which garden-path the hearer and require conscious reanalysis, description grammar must be combined with a preference system that comes up with a specific reading corresponding to the hearer's currently preferred interpretation.

LDG presupposes an embedding of linguistic analysis in a model of human reasoning in general. This seems to be appropriate, since in determining the meaning of a sentence in its context of utterance, a language user employs sources of information other than just his grammatical knowledge and the linguistic input. To disambiguate a sentence, or to determine how a sentence relates to its discourse context, he makes use of his world knowledge and reasons about the nonlinguistic context of utterance as well. Since parsing is modeled as reasoning in LDG, the language user's access to common knowledge and other sources of information in the interpretation process can naturally be integrated in the model of sentence interpretation.

Finally, as a constraint-based system LDG is an open, flexible model of linguistic analysis. The grammar is nothing but a set of constraints formulated in one and the same logical language, simultaneously constraining the properties of linguistic structures. As such, the interaction of constraints coming from different grammatical levels (such as syntax and semantics) is modeled straightforwardly, and in principle new constraints, instantiating other grammatical levels, can simply be added. It also means that LDG is very useful as a logical toolkit for linguists. No particular linguistic theory is hard-wired in the description grammar: the formalism provides the means for the specification of a linguistic theory, rather than that it implements a particular linguistic theory itself. Thus, the formalism can very well be employed to implement and test alternative constraints and to compare different linguistic theories.

Chapter 3

LDG for Discourse Analysis

3.1 Introduction

The LDG presented in the preceding chapter is a sentence grammar, but the formalism can be employed for the analysis of discourse as well. In this chapter, we put forward a framework of discourse interpretation which incorporates an incremental variant of LDG. Beside the grammar system, the framework includes a reservoir of general or world knowledge, and a default inference component. As before, the grammar system generates logical descriptions of linguistic structures. The descriptions represent the language user's knowledge of the ongoing discourse. Given his grammar and his world knowledge, the language user may reason about the tree models verifying the description of a discourse. The tree models figure as discourse representations. Each different tree model stands for a possible linguistic analysis, and correspondingly, a reading, of the discourse.

The integration of LDG in a model of discourse interpretation is interesting for various reasons. Firstly, due to its formulation in terms of description theory, LDG provides an elegant and powerful treatment of underspecification and ambiguity at discourse level. As can be seen from the work available on the application of underspecifiation theory to linguistic phenomena exceeding sentence level this is a wide and fruitful field of research indeed.

To name just some of this work, building on Reyle's UDRT (1993), Asher and Lascarides (2003) propose underspecified logical forms in SDRT for the treatment of various forms of ambiguity, such as scope ambiguity of quantifiers, lexical ambiguity, anaphora resolution, VP-ellipsis, presupposition binding or accommodation, structural ambiguity (attachment ambiguity) of discourse connectives and relations of coherence, and semantic underspecification of the latter. Furthermore, the Saarbrücken school of description grammar (Pinkal 1996; Niehren et al. 1997; Egg et al. 1998) focuses on the treatment of quantifier scope, VP-ellipsis, parallellism phenomena and anaphora resolution. Blackburn and Bos (1999) provide a

computational treatment of anaphora and presupposition resolution. Finally, in Gardent and Webber (1998), Blackburn and Gardent (1998), Duchier and Gardent (1999), Duchier and Gardent (2001), and most recently Egg and Redeker (2006) various forms of ambiguity and underspecification of coherence or discourse relations are investigated. A less well investigated, but similarly interesting topic of research would be the treatment of focus articulation in English in an underspecifiation formalism. We build on all of the above work in developing our own model of discourse interpretation.

Another benefit of using LDG in a model of discourse interpretation is that the resulting model is open to the testing of cognitive claims. LDG models parsing as reasoning. Employing grammatical knowledge, world knowledge, and specific knowledge of the sentence or discourse under consideration, a language user reasons about the linguistic analysis and meaning of a sentence or discourse. As part of a general model of human knowledge and reasoning, our framework of discourse interpretation naturally provides a cognitive theory of discourse analysis. Thus, predictions of the framework can be tested on their cognitive tenability and attractiveness.

Finally, again due to its formulation in terms of description theory, LDG functions as a *specification language* for discourse theory. LDG allows us to 'talk' about the properties of a discourse theory in a logical language. Of course, in what follows we will make certain choices and fill in the specifics of some constraints of the discourse theory. However, for the particular choices one can make in discourse theory there are often viable alternatives. In our view it is a step forward to be able to implement and consider alternative treatments in one and the same formalism. Different approaches to a given issue can then be compared, and maybe in some points a synthesis can be obtained. Our framework of discourse interpretation provides a platform for this.

In this chapter the basic formal characteristics of the LDG framework of discourse interpretation are presented. The focus is on discourse syntax and the effects of incrementality in reasoning about discourse descriptions. An explanation of the semantic representation language employed in the framework and the semantics of discourse contributions will be given in the next chapter. The same goes for central notions such as discourse meaning, felicity or appropriateness conditions, and the integration of local contexts in LDG.

In section 3.2 the framework of discourse interpretation is introduced and its overall working is explained. In section 3.3 we zoom in on the relation between sentence grammar and discourse grammar. It is argued that sentence theory is a fully integrated part of discourse theory in our model. In section 3.4 we talk about discourse structure. The grammar is extended to cover syntactic structure above clause and sentence level. Our use of the terms 'discourse constituent', 'discourse relation' and 'discourse operators' is explained, and a simple treatment of discourse operators and discourse relations in the lexicon of the grammar is

provided.

In section 3.5 we present an incremental version of the LDG formalism of Muskens (2001), which was discussed in the preceding chapter. The implications of incrementality for the relation between input descriptions, discourse descriptions and the tree structure models that verify them are discussed (section 3.5.1). Input descriptions are adapted to the incremental set-up of the grammar (section 3.5.2), and the incrementation operation is defined (section 3.5.3). The resulting LDG formalism is applied to an example (section 3.5.4). The example features two sentences linked by an implicit discourse relation. Given the properties of the grammar, it can be deduced that there is a discourse relation linking the tree representations of the two sentences. The relation is semantically underspecified. The hearer is predicted to infer an appropriate semantics, however, employing his world knowledge and default reasoning.

In two subsequent sections, some further properties of incrementation as it is modeled in our framework are discussed. Section 3.5.5 defines structural accessibility of the nodes named in a discourse description, and it explains what it means for an input description to attach at the right frontier of a described discourse tree. Section 3.5.6 discusses the deterministic hypothesis, a central cognitive claim which applies to the attachment of new input descriptions in incremental LDG. The hypothesis says that cases of attachment ambiguity which require no conscious reanalysis of the hearer are processed deterministically. The implications of the hypothesis for the framework of discourse interpretation are briefly discussed.

In section 3.6 the chapter is summarized and some loose ends are indicated.

3.2 A Framework of Discourse Interpretation

We propose a framework of discourse interpretation that incorporates an LDG system which is adapted to cover incremental discourse processing. In doing so we implement what might be called the central idea of discourse grammar theory, namely that discourse, like sentences, is subject to rules of syntactic and semantic composition—that, in other words, the notion of a *grammatical* discourse makes sense.

Discourse grammaticality is a more encompassing notion than sentence grammaticality, however. The object of study is not the syntax and semantics of sentences in isolation, but rather the relation of sentences (or utterances, or discourse contributions) to their context of interpretation. For a discourse to be grammatical, the sentences uttered in it must not only be syntactically correct but also semantically well-formed in their context of interpretation, so that the discourse as a whole is coherent.

The framework of discourse interpretation that we propose takes the interaction of the various factors that make a discourse into a coherent whole as its point of departure. The grammar is a set of constraints on the syntactic and semantic

properties of discourse. A grammatical discourse satisfies all these constraints. Different types of discourse contributions, such as assertions, denials, confirmations, objections, etc., can be characterised in terms of the particular constraints which make them felicitous in their context of interpretation as contributions of that type.[1] In actual discourses, it may occasionally be unclear what the type of a given discourse contribution is. The grammar formalism predicts this, as it allows the relation of a discourse contribution relative to its context of interpretation to be underspecified.

There is more to discourse interpretation than just the application of grammatical knowledge, however. It is commonly assumed that hearers employ world knowledge in, for instance, testing presuppositions, resolving bridging anaphora, and inferring implicit coherence relations. Furthermore, in interpreting ambiguous or underspecified discourse they often have preferences for certains readings over others. Hence we include both world knowledge and a default reasoning component generating preferences in the framework of discourse interpretation.

A *discourse model* or *framework of discourse interpretation* of a language user a, $\mathcal{F}_a = \langle G^+, \mathcal{K}, \blacksquare \rangle$, consists of an incremental grammar system $G^+ = \langle \mathcal{G}, \Delta_0, I, + \rangle$, a reservoir of general background information or world knowledge $\mathcal{K}$, and a default reasoning component $\blacksquare$. The incremental grammar system combines a discourse grammar $\mathcal{G}$ with *discourse descriptions* that are build up incrementally from the initial discourse description Δ_0 and input descriptions in I by means of the operation $+$. As before, $\mathcal{G}$ represents a's grammatical knowledge of a given natural language. $\mathcal{K}$ represents his nonlinguistic knowledge. This includes world knowledge, common sense knowledge, and knowledge about the actual utterance situation. The inference engine $\blacksquare$ stands for his default reasoning capacity and his capacity to assign 'preferences'. As in the preceding chapter, standard logical deduction is available to the language user as well, of course.

$\mathcal{G}$ is a Logical Description Grammar. As before, it is a set of *descriptions*, or statements in classical type logic. For any given stretch of discourse the grammar system produces a *discourse description*. This is a set of logical statements describing the syntactic and semantic relations characterising the discourse processed so far. Discourse descriptions represent the language user's knowledge of the discourse. As before, discourse descriptions can be partial; they figure as underspecified discourse representations in the discourse model.

A language user may reason about the *models* of a discourse description given $\mathcal{G}$ and $\mathcal{K}$, employing logical deduction and $\blacksquare$. As before, the verifying models of

[1] Discourse contributions can be characterised in our discourse model to the extent that the parameters made available in it can be set for each type of discourse contribution. Parameters that are currently not covered in the discourse model, for example the plans and intentions of the participants in making a particular kind of discourse contribution, or their attitudes and commitments, do not feature in our formal treatment of discourse contributions. In this respect, the model offers only a partial characterisation. For some discussion of integrating participants' commitments in the discourse model, see section 3.2 of the next chapter.

a description are tree structures. They are comparable to the 'discourse parse tree' in the Linguistic Discourse Model (Polanyi 1988; Scha and Polanyi 1988) and (L)TAG for discourse (Webber and Joshi 1998; Webber et al. 1999). The tree structures figure as fully specified representations of the discourse processed so far. As before, aspects of the syntax or semantics of the discourse can be underspecified, in which case there will be more than one verifying tree model.

The default reasoning component ■ stands for the language user's capacity to reason about world knowledge in interpreting discourse, and to make choices in the face of uncertain or incomplete information. We assume the language user is able to weigh the available information and to decide whether some options are for instance most likely to be true, or most uncontroversial, or most desirable. On the basis of this weighing he is able to order alternative verifying models according to preference, and to opt for the most preferred ones. Crucially, conclusions drawn by means of ■ are defeasible, and can be retracted when more information becomes available to the language user. Unlike the constraints in $\mathcal{G}$ the preferences generated by ■ are 'weak constraints'.

Although we will not say anything about the internal working of ■, its presence will help to delineate the tasks within the framework of discourse interpretation, and clarify the impact of default reasoning on the disambiguation of discourse. As a component in the framework, ■ is comparable to what used to be called DICE, and is now part of the logic of information packaging in SDRT (Asher and Lascarides 2003), to the general inference mechanism for defausible aspects of discourse semantics proposed in Webber et al. (1999), and to the system of preferences or biases presupposed in Gardent and Webber (1998). As we envisage its function and working in $\mathcal{F}_a$, however, it most closely reflects the type of reasoning and the weighing of alternatives proposed in theories of plausible reasoning (Rescher 1976) and belief change (Gärdenfors 1988).

Abstracting over individual language users, a logical description grammar for discourse analysis $\mathcal{G}$ may be taken to characterize a natural language. It provides an analysis of all possible 'grammatical discourses' in a language. As we said before, we are using the term grammatical in a wide sense. For a discourse to be *grammatical*, the sentences uttered in it must be syntactically and semantically well-formed in their context of interpretation. A description grammar for discourse $\mathcal{G}$ is *adequate* for a given language if, for any grammatical discourse in that language and for no ungrammatical one, the description of the discourse is consistent with $\mathcal{G}$. A discourse model $\mathcal{F}_a$ is adequate if it generates the (cognitively) correct preferences over the verifying models of discourse descriptions given $\mathcal{G}$, and $\mathcal{G}$ is adequate.

We define a *discourse* as a sequence of utterances $u_1 \ldots u_n$, where the linguistic expressions uttered are sentences. Utterances may be either written or spoken. In the case of spoken discourse, we make the simplifying assumption that the utterances are strictly ordered according to onset time. Furthermore, we assume that each utterance is the vehicle of a single discourse contribution. A *discourse contri-*

bution is a 'dialogue move' or 'language act' made by a particular participant at a certain point in a given discourse, with a particular rhetorical function relative to the discourse context, for example an assertion, a denial, a confirmation, or a question.

Unlike the LDG formalism presented in the preceding chapter, G^+ processes discourse *incrementally*. Every discourse contribution induces an incrementation of the discourse description built up so far. Incrementations are always monotonic, even when the discourse contribution causing it is a correction or a denial. An incrementation step never leads to the retraction of a statement from the discourse description, it only leads to a further specification of the tree model of the discourse under consideration. Only when the language user, confronted with new information, is forced to cancel default conclusions drawn earlier by means of ■, real backtracking takes place.

The following schema shows a model of incremental processing in our discourse model. In principle, u_i may stand for any relevant processing unit, but since we are focusing on discourse level we assume they stand for utterances of sentences.

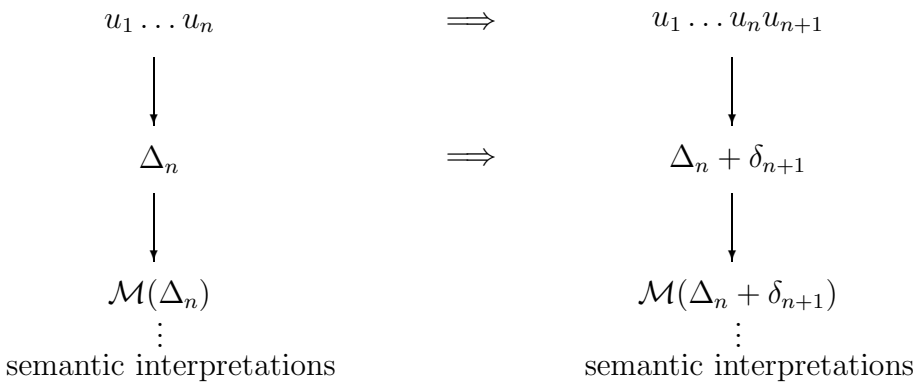

Figure 3.1: An incrementation step in LDG for discourse

For any utterance u_i, δ_i is the *input description* resulting from analysing or parsing the sentence uttered. $u_1 \ldots u_n$ is an initial stretch of discourse, Δ_n the discourse description resulting from processing it, and $\Delta_n + \delta_{n+1}$ denotes the result of incrementing Δ_n with δ_{n+1}. Discourse descriptions are basically sets of input descriptions (but see section 3.5.3). An incrementation step ($\Longrightarrow$) consists in processing a new utterance u_{n+1}; the grammar then produces an input description δ_{n+1} of u_{n+1} and combines this with the current discourse description Δ_n, so that a new discourse description $\Delta_n + \delta_{n+1} = \Delta_{n+1}$ is formed. Furthermore, $\mathcal{M}$ is a function which produces the set of models or trees which satisfy a discourse description, given $\mathcal{G}$. Each verifying model in $\mathcal{M}(\Delta_n)$ or $\mathcal{M}(\Delta_n + \delta_{n+1})$ comes with a semantic interpretation, representing a possible reading of the discourse.

We will use the term *update* to refer to the result of incrementing a discourse description by an input or sentence description at a particular moment in processing the discourse; this is represented by $\Delta_n + \delta_{n+1}$ in the schema. At the level of the verifying tree models, updates are characterised not just in terms of the tree structures described by, respectively, a given input description and a given discourse description, but also by the attachment possibilities of the two. Hence, an update is essentially more than just an input description. While input descriptions represent the *sentences* uttered in a discourse, updates represent the *discourse contributions* made through these utterances.

3.3 Sentence and Discourse Theory Integrated

An *incremental LDG system* or *LDG for discourse analysis* $G^+ = \langle \mathcal{G}, \Delta_0, I, + \rangle$ consists of the following components.

- A grammar $\mathcal{G} = \langle \mathcal{A}, \mathcal{W} \rangle$, where $\mathcal{A}$ is a set of axioms or general descriptions and $\mathcal{W}$ a set of lexical descriptions (classifying and elementary),
- I, an infinite set of input descriptions δ, generated per processed utterance,
- Δ_0, the initial discourse description, and
- $+$, the incrementation operation.

Apart from the last two, these components are familiar from the grammar system presented in the preceding chapter. As before, $\mathcal{G}$ is a Logical Description Grammar. The general descriptions or axioms $\mathcal{A}$ regulate the types of structures that can be described by a discourse description generated by the grammar. Though some of the axioms will be modified with a view to the treatment of implicit discourse relations and the change to an incremental set-up, and some general descriptions will be added, their general purpose and effect is the same as before. The lexical descriptions in $\mathcal{W}$ characterise the elementary tree structures contributed by lexemes of the different word classes; together they constitute the lexicon of the grammar. As before, input descriptions name the lexemes occurring in the linguistic expression uttered and state the linear order of these lexemes. Unlike the input descriptions in the original proposal, however, they do not contain a circumscription clause.[2] The new components in the list embodies the change from a nonincremental to an incremental grammar formalism: $+$ constructs a new discourse description from a given discourse description plus a new input description, Δ_0 is the start element of this.

In the preceding chapter, the term *linguistic theory* was used to refer to any particular theory of sentence analysis laid down in $\mathcal{G}$. In this chapter, we prefer

[2]See section 2.4.5 of the preceding chapter.

to use the term *discourse theory* instead, and this then refers to any particular theory of discourse analysis laid down in the framework, or discourse model $\mathcal{F}_a$ as a whole. There is no specific component within the discourse model which represents the discourse theory as such. Different components of the grammar cover different aspects of the discourse theory. For example, the assumption that discourse structures are tree structures is laid down in the axioms $\mathcal{A}$, and the syntax and semantics of discourse connectives is laid down in the lexicon $\mathcal{W}$. As we shall see in subsequent chapters, the resolution of anaphora and the satisfaction of presuppositions is governed by constraints in both $\mathcal{W}$ and $\mathcal{A}$. The incrementation operation + captures another property of the discourse theory. Finally, the impact of default reasoning and preferences on the disambiguation of discourse in general is given by ■.

The discourse theory laid down in the discourse model comprises both sentence analysis and discourse analysis. As with LDG for sentence analysis, the lexicon is the motor of the syntactic and semantic analysis of the linguistic input. Indeed, if we consider an LDG for sentence analysis and an LDG for discourse analysis characterising one and the same natural language, the lexicon of the second is a strict extension of the lexicon of the first. Thus, sentence theory is a fully integrated part of discourse theory.

Notice that this is not in any way prohibiting the implementation of a notion of 'sentencehood' in the grammar, as a particular syntactic domain over which constraints can be stated. In fact, we believe sentencehood is important in controlling unwanted ambiguity and the attachment possibilities of subsequent discourse. Even when a notion of sentencehood is not explicitly implemented in the formalism, however, the term *sentence theory* can sensibly be employed to refer to sentence internal relations, or to linguistic constraints in the grammar that are specifically operative at, or below the level of sentences. There are, however, many constraints in the grammar, for instance on anaphora resolution, which apply across the board, so that it is sometimes a matter of perspective rather than strict delineation that makes us talk about sentence theory rather than discourse theory.

In computational approaches to discourse analysis it is common to distinguish between sentence parsing and discourse parsing, see for example Polanyi et al. (2004), or Webber (2004) and Forbes et al. (2001). In talking about discourse processing, we find it practical to make this distinction as well. We adopt an idealised picture of discourse processing, in which the utterances in the discourse are processed one after another and there is a component in $\mathcal{G}$ which delivers the appropriate input descriptions for each utterance, independently of the discourse processed sofar. Utterances of sentences are the standard processing units of the discourse grammar. This can be changed whenever it is desirable, however. There is no principled reason why we should not enter the domain of sentence analysis and, for example, in the processing of a complex sentence treat clauses or simple sentences as processing units.

In general we expect that there are subtle interactions between sentence analysis and discourse analysis. For example, the topical structure of a discourse raises expectations about the form and content of the next utterance to be processed. Given the fact that the sentence grammar and the discourse grammar are integrated in a single formalism we are hopeful that such interactions can succesfully be modeled in the discourse model.

3.4 Discourse Structure and Discourse Relations

All theories of discourse interpretation claim that discourse is *structured*, and that the structure of the discourse processed so far constrains the form and interpretation of incoming utterances. Discourse structure results from the sequential order of the utterances in a conversation and the rhetorical links between them.

The type of hierarchical structure proposed varies per theory. The most influential approaches are, on the one hand, Mann and Thompson (1988)'s Rhetorical Structure Theory (RST), and on the other, strict tree structure models such as Polanyi's Linguistic Discourse Model and Webber's LTAG for discourse. RST assumes a hierarchical structure which is built up from rhetorical relations linking pairs of nonoverlapping text spans; one of the pair is the *nucleus* and the other the *satellite*. A single nucleus can be associated with several satellites, each via their own rhetorical relation. Hence RST structures are more flexible than strict tree structures. In our model we opt for the simplest type of structure. Following Polanyi's Linguistic Discourse Model and LTAG for discourse, we will assume that unambiguous discourse can be represented by single tree structures, *discourse (parse) trees* composed of discourse constituents linked by discourse relations.

Various arguments in favour of the idea that discourse is hierarchically structured through rhetorical links are given in Asher and Lascarides (2003), chapter 1. Most of this concerns the resolution of anaphoric and presuppositional expressions. Generally speaking, if an incoming utterance refers to individuals, events or concepts introduced earlier in the discourse, it depends, among other things, on the structural position of the antecedent whether the link can be expressed through personal pronouns, definite descriptions or other types of anaphors.

Let us take a quick look at one of these arguments. Consider the following discourse, discussed in Asher and Lascarides (2003) pp. 8–10.

(3.1)a. *Max had a great evening last night.*
b. *He had a great meal.*
c. *He ate salmon.*
d. *He devoured lots of cheese.*
e. *He then won a dancing competition.*
f. *# It was a beautiful pink.*

The last utterance in the discourse is a statement about the salmon Max ate, introduced earlier in (3.1c). Intuitively speaking, (3.1f) elaborates on (3.1c), but the proposition expressed by (3.1c) is somehow out of reach. In particular, we have trouble resolving the pronoun 'it' to the salmon Max ate. In discourse grammatical approaches, this can be explained on the assumption that the discourse unit containing the antecedent of a personal pronoun in the discourse representation must be in a particular structural position relative to the discourse unit representing the clause containing the pronoun: it must be at the 'right frontier' of the discourse tree. To see this, consider the following picture, which is a tree representation of the part of this discourse consisting of the utterances (3.1a–e). The labels *elaboration* and *narration* denote coherence, or discourse relations.[3]

(3.2)

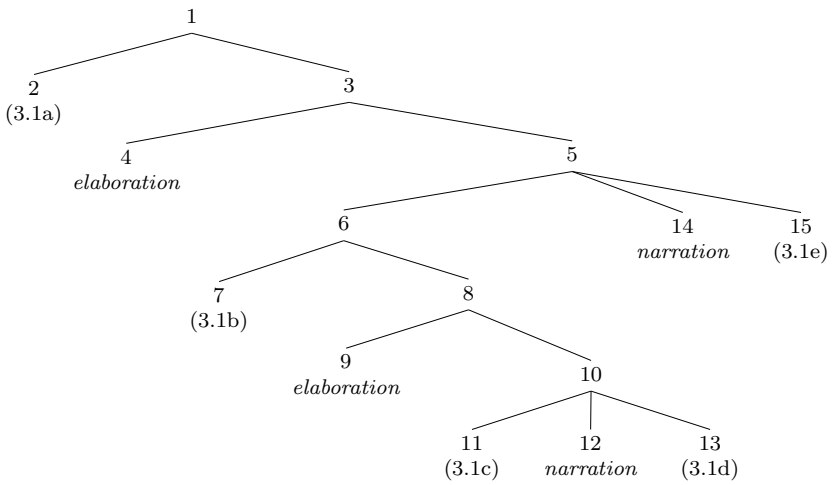

The right frontier of a discourse parse tree in LTAG for discourse and the Linguistic Discourse Model is roughly the set of nodes in the tree representing the

[3]The tree structure in (3.2) uses the discourse relations occurring in the SDRT representation of this example in Asher and Lascarides (2003) on page 140, but it presupposes the discourse syntax of our own grammar. We constructed the tree structure in (3.2) by projecting the π-labels in the SDRT representation into nodes or discourse units and translating the SDRT subordination structure as syntactic dominance. (π-labels are discourse referents for speech acts in SDRT.) The hierarchical structure of SDRSs, however, conforms to Rhetorical Structure Theory, and differs essentially from our discourse syntactic structures in that it permits a single π-label to serve as an argument in more than one discourse relation. To allow this in our grammar too, we would need graphs or tree structures of a more flexible kind, in which a node can have more than one mother node. See Blackburn and Gardent (1998) and Duchier and Gardent (1999) for some discussion of these so-called *re-entrant* trees. (In fact, as noted in Sporleder and Lascarides (2007), SDRT goes a step further than RST in that it allows for the possibility of more than one relation holding *simultaneously* between two discourse units.) On the 'treeness' of discourse structure, see also Egg and Redeker (2006). While adopting strict tree structures, they discuss some constructs in discourse that are often claimed to require more flexible structures, namely 'crossed dependencies', discontinuous (that is, interrupted) constituents, and discourse units that enter into more than one coherence relation with other discourse units. Especially the last type of construct poses a challenge to plain tree structure in their view.

smallest rightmost discourse unit and the discourse units dominating it. In (3.2) it consists of the nodes 1, 3, 5, and 15. As an effect of the incrementality of the discourse parsing process, tree structures representing incoming utterances attach only at the right frontier of the discourse tree. The antecedent of the pronoun in (3.1f) is introduced at node 11, but this is not a discourse unit at the right frontier of the discourse tree. Consequently, the pronoun cannot refer to it. Whatever the attachment position of the tree representation of (3.1f) in (3.2) is going to be, the use of a pronoun is predicted to be infelicitous.

Observe that, although the use of a pronoun at this point in the discourse is inappropriate, it is not impossible to refer to the salmon Max ate by other means. The speaker might say something like, 'By the way, the salmon Max ate was a beautiful pink.' Discourse grammatical theories will say that the definite description, unlike the pronoun, is sufficiently informative to re-establish the salmon as a topical element and to resolve or accommodate the presupposition that there is some salmon eaten by Max, given the attachment of the tree representation of the incoming utterance at a suitable node on the right frontier of the discourse tree structure.[4] Presumably, this is node 5, where Max's having a great evening last night is the topic under discussion.

We will come back to this example in chapter 5 when we discuss the implementation of discourse structural constraints on anaphora resolution in LDG. In the rest of this section we will talk about the building blocks of discourse structure, namely discourse units, discourse relations and discourse operators. We want to make clear what these notions stand for in the context of our LDG for discourse. We will also discuss how implicit, or inferred discourse relations contribute to discourse syntax in LDG. For definitions of attachment, structural accessibility and the right frontier in LDG we refer to section 3.5.5.

Discourse units and discourse operators As we said, we presuppose that an unambiguous discourse can be represented by a single tree structure, a discourse tree composed of discourse constituents linked by discourse relations. Thus, as in LDG for sentence analysis, the discourse grammar generates descriptions of tree structures. The notion of a *discourse unit* is somewhat fluid, due to the fact that the 'discourse unitness' of clauses is established a posteriori. That is, if in a subsequent update a discourse relation attaches to a clause, we know it is a discourse unit. Syntactically, discourse units may loosely be defined as clauses, or constituents composed of clauses and discourse operators. We leave open whether every clause can serve as a discourse unit.[5] All clauses and discourse units in

[4]The discourse cue phrase 'By the way' plays a role in this too. It signals a change from the latest local topic, namely 'Max won a dancing competition'.

[5]But see Polanyi et al. (2004) on this issue. Compare also Marcu (2000) on the problem of rhetorical grounding, that is, the task of automatically determining the elementary units of text in discourse parsing. Marcu speaks of 'clause-like units'.

our grammar are of the syntactic category S.[6] Semantically, discourse units are composed of clause meanings, and meanings of discourse operators and discourse relations.

We use the term *discourse operator* to refer indiscriminately to coordinating and subordinating connectives ('and', 'but', 'or', 'if','because', 'unless' etc.), sentence adverbials or modifiers ('furthermore', 'however', 'then', 'in addition', 'on the other hand', etc.), discourse cues ('OK', 'well', 'so' and 'anyway', etc.) and the discourse particles 'Yes' and 'No'. The lexicon of the grammar contains lexical descriptions for lexemes of these word classes. Like all other lexemes, discourse operators come with a particular truth-conditional contribution to the semantics of the discourse. As such they constrain the discourse relations which make a discourse into a coherent whole. They do not necessarily fully determine them, however.[7]

Here are graphical representations of the lexical descriptions of a few of them. For the time being, we suppress the semantic values in these pictures; they will come in a variant of compositional DRT which will only be explained in the next chapter. The same goes for the felicity conditions or context-sensitive constraints that are often attached to individual discourse operators.

(3.3)

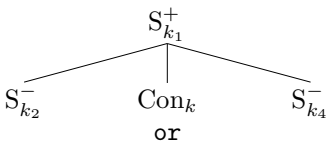

The coordinating connective 'or' takes a clause or discourse unit as an argument to its left and another one as an argument to its right. We will sometimes refer to S_{k_2} as its *contextual argument* and to S_{k_4} as its *internal argument*. Semantically, 'or' results in a disjunction of the semantic values of its arguments.

[6] Alternatively, a distinction between S-nodes, for clauses, and T or D_c nodes for 'texts' or 'discourse units' could be introduced. This is the approach taken in, among others, Webber (2004)—and which, incidentally, was followed in Van Leusen and Muskens (2003). Webber allows clauses or sentences, which can function as minimal discourse units, to anchor D_c nodes. A sentence level LTAG parser delivers the input structures for a discourse level D-LTAG parser. Since many discourse operators function both below and above sentence level, this set-up of the grammar implies that there is a systematic ambiguity between 'sentence' and 'discourse' uses of these operators. This is perfectly acceptable, but it does raise the question whether in all these cases there is indeed a corresponding difference between the semantic-pragmatic impact of the operator at discourse level and its impact at or below sentence level.

[7] Many discourse operators are known to be ambiguous, or compatible with more than a single discourse relation or coherence relation. See for instance Hirschberg and Litman (1993) or Knott and Mellish (1996).

(3.4)

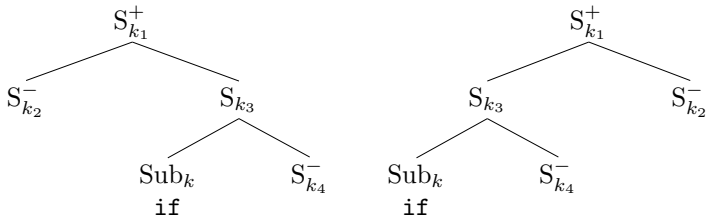

The subordinating connective 'if' takes an argument to its left, this is the subordinated clause S_{k_4}, and an argument which may either precede or succeed it, the main clause S_{k_2}. The lexical description of 'if' underspecifies the precedence order of the main and the subordinated clause; in (3.4) this is expressed by showing both options. Semantically, 'if' results in a conditional statement.

(3.5)

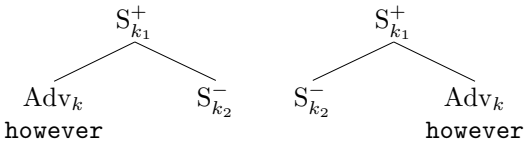

The lexeme 'however' is an adverbial, which may either precede or succeed its internal argument, that is, the clause or discourse unit it modifies. Again, the lexical description underspecifies this. Semantically, 'however' reproduces the semantic value of its internal argument, and triggers a requirement of contrastivity or denial of expectation relative to its context of interpretation.[8]

[8]Alternatively, discourse adverbials like 'however' could be analysed as two-place operators, as in

(3.6)

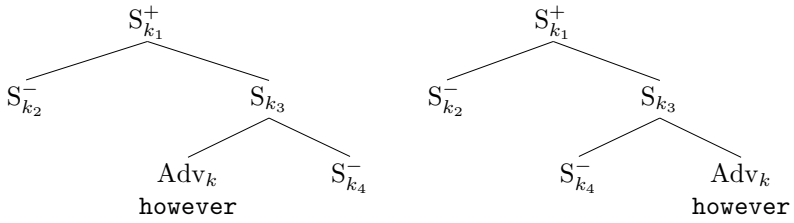

This would imply that the relation to the discourse context expressed by the adverbial is a *structural* one, rather than a presuppositional or anaphoric one, as it will be in our account. An advantage of the simpler structure in (3.5) is that potential problems in the syntactic composition of sentences containing sequences of sentence or clause operators (as in, for example, "John promised to come. If, however, against all expectations, he doesn't, one of us will have to give the talk.") can be avoided. Some further reasons for preferring the structure in (3.5) over the one in (3.6) are given in Webber et al. (2003). Webber c.s. argue that discourse adverbials behave anaphorically in that they admit the crossing of structural dependencies, and have a wider range of options with respect to their external argument than do coordinating and subordinating connectives.

(3.7)

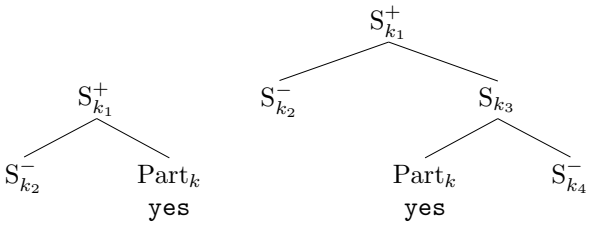

Finally, 'yes' is a discourse particle. It may occur on its own, operating on a single argument to its left, its contextual argument. It may also come with an additional argument to its right, as in 'John is a nice guy. Yes, he certainly is'. Semantically, it produces the semantic value of its contextual argument, or the combination of this with the semantic value of its internal argument. In both cases the particle serves to express the confirmation of its contextual argument. Additional constraints on the context of interpretation regulate this.

The lexical descriptions just presented merely serve to show a possible treatment in LDG of these discourse operators. Different choices could be made, and further research into the meaning and use of these operators in their discourse context will certainly lead to the exploration of alternative treatments.

Discourse relations and implicit anchors Literature on discourse analysis abounds in names and classifications of rhetorical relations, coherence relations, and discourse relations, compare for instance Hobbs (1985), Polanyi (1985), Mann and Thompson (1987), Grosz and Sidner (1986), Webber (1991), Knott and Mellish (1996), Marcu (1997), and Asher and Lascarides (2003). The distinctions between these three kinds of relations are not always drawn very sharply. Sometimes the terms are interchangeable, and sometimes they overlap. We shall use the term 'rhetorical relation' only when talking about speech acts or discourse contributions, and the particular semantic-pragmatic relations to the conversational context that characterise them as contributions of a certain type. We shall use the term 'coherence relation' in a more general sense, to refer to the semantic-pragmatic relations that glue together the utterances in a discourse. This covers relations of a rather 'rhetoric' flavour such as 'correction' and 'confirmation' as well as what might be called content linking relations, such as 'contrast' and 'parallel'. Finally, we will employ the term 'discourse relations' to refer to the elements in our discourse grammar that denote coherence relations.[9]

We will employ the same names for discourse relations as Asher and Lascarides (2003) do, and in our syntactic-semantic characterisation of them try to keep as close to theirs as we can, given the possibilities of our grammar formalism. We will freely use terms like 'correction', 'counterevidence', 'elaboration', 'narration', 'contrast', 'explanation', and 'result', trusting that in the context in which we are

[9]This use of the term 'discourse relation' is specific to our approach. In most approaches, it refers simultaneously to the element in the grammar or discourse theory and the (discourse / coherence / rhetorical) relation it represents.

using these terms it is intuitively clear what they refer to. In the next chapter we will spell out the semantics of some of these relations in LDG.

Discourse relations, in our grammar, are implicit elements denoting coherence relations. Syntactically, they link clauses to their discourse context. Semantically, they characterise coherence relations through the particular combination of their compositional semantics and a (possibly empty) set of context sensitive constraints or *felicity conditions*. Felicity conditions constrain the relation between the semantic values of the arguments of a discourse relation in the relevant context of interpretation or local context. As an integrity constraint of the grammar we require that it describes only a finite number of discourse relations, that is, discourse relations are a closed class. With a view to a classification of discourse relations, additional integrity constraints on the relation between the defining properties of discourse relations and the semantics of discourse operators might be added.

Coherence relations may be completely determined through the presence of explicit discourse operators and the tense, mood and modality of their internal argument. They are often underspecified, however. Notably, this is the case when no overt discourse operator is present, and the relation of coherence must be inferred by the hearer on the basis of his grammatical knowledge, world knowledge, and information about the utterance situation. The following example illustrates this.

(3.8) *John left early.*
He was ill.

In interpreting the discourse, the hearer infers that there is some relation between the fact that John was ill and the fact that he left early. In all probability the hearer assumes that John left early because he was ill, that is, the second sentence provides an explanation of the proposition expressed by the first. The relation of explanation is what makes the discourse into a coherent whole. In this situation we assume that there is an implicit linguistic element representing the coherence relation: a discourse relation.

Like many other authors, for example Lascarides and Asher (1991), Lascarides et al. (1992), Webber and Joshi (1998), we assume that discourse relations, like overt discourse operators, contribute to the syntactic structure of discourse. Thus, for the discourse in (3.8), given the interpretation suggested, our LDG will generate a description of the following tree structure.[10]

[10]Other interpretations or readings of this discourse are possible. The tree structure in (3.9) corresponds to the most preferred reading or interpretation of the discourse. In section 3.5.4 we say more about this.

(3.9)

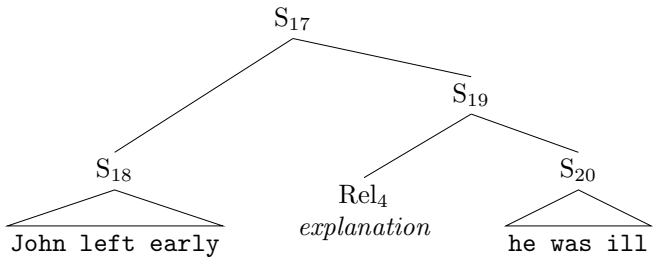

In line with Webber and Joshi (1998) and Webber et al. (1999), we assume that discourse relations can anchor elementary structures. The lexicon of the grammar contains two elementary descriptions for elements of word class *rel*: one for coordinating discourse relations, such as narration and contrast, and one for subordinating discourse relations, such as elaboration, explanation, result and correction. The following pictures of descriptions feature these two elementary forms. Again, for the semantics, see chapter 4.

(3.10) *rel:*

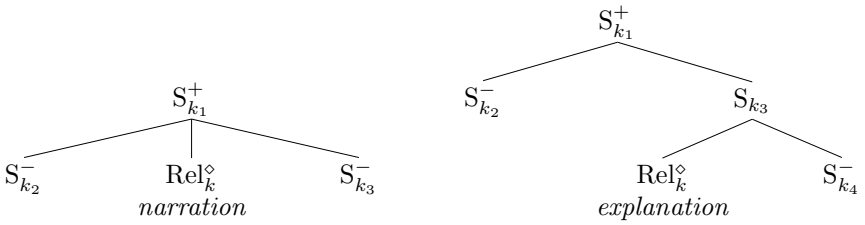

As in the preceding chapter, the sign $^\diamond$ marks the anchor of the elementary structure. In the case of coordinating elementary structures, k negatively anchors itself, k_2 and k_3. It positively anchors itself and k_1. Coordinating discourse relations require a discourse unit S_{k_2} to their left (the contextual argument) and a discourse unit S_{k_3} to their right (the internal argument). In the case of subordinating elementary structures, k negatively anchors itself, k_2, k_3, and k_4. It positively anchors itself, k_3, and k_1. Subordinating discourse relations require a discourse unit S_{k_2} to their left (the contextual argument) and a discourse unit S_{k_4} to their right (the internal argument).[11] In the discourse tree in (3.9) the nodes 17, 18, 19, 20 and 4 instantiate the elementary description of subordinating discourse relations, selected by the discourse relation *explanation*.

[11] Various alternative treatments might be considered. For instance, one might opt for a single elementary structure for all discourse relations, and treat the distinction between coordinating and subordinating relations as an independent syntactic-semantic property. This is more in line with for example the approach of Asher and Lascarides (2003), where the 'subordination relation' functions as a defining property of certain discourse relations, constraining their impact on the discourse topical structure. Another alternative one might consider is to adopt a more flexible elementary structure for subordinating relations, or possibly for *all* discourse relations, of the following form:

In order to allow for nodes and structure in a discourse tree that are not anchored to the lexemes occurring in the discourse, but to discourse relations, some slight modifications of the axioms which regulate lexicalisation in LDG are necessary. Remember that LDG adheres to the lexicalisation principle adapted from LTAG: all nodes in a verifying tree model of a description must be positively and negatively anchored to a lexical node in the tree. In the case of discourse relations, however, there is no lexeme which anchors the nodes and structure introduced in the discourse description. We account for this by introducing a distinction between *implicit anchors* and *lexically expressed* ones in the grammar. Discourse relations come with the first type of anchor, and lexemes with the second type of anchor.

In the LDG formalism presented in the preceding chapter, nodes which serve as an anchor are lexical nodes (and vice versa), and all lexical nodes carry an explicitly expressed lexeme. Moreover, terminal nodes and lexical nodes coincide. In order to cope with implicit anchors we differentiate between lexically expressed nodes and terminal nodes. *Terminal nodes* are nodes which do not dominate any other node:

$\mathcal{A}15\ \ \forall k[\, term(k) \ \leftrightarrow\ \neg\exists k' \colon k \lhd^{+} k']$

Anchors are now required to be terminal, rather than lexical. We adjust the anchoring axioms $\mathcal{A}12$ and $\mathcal{A}13$, replacing '*lex*' with '*term*'.

$\mathcal{A}12'\ \ \forall k : term(\alpha^{+}(k))$

$\mathcal{A}13'\ \ \forall k[\, k = r\ \vee\ term(\alpha^{-}(k))]$

Anchors can be either lexical or not, and in the latter case we call them implicit anchors. 'Lexical' now means lexically expressed, or not phonologically empty. We introduce the predicate *imp* to indicate that a node is terminal but not lexical. This makes the node into an *implicit anchor*:

(3.11)

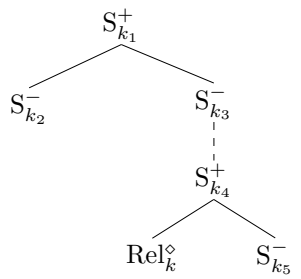

In comparison with the structure in (3.10), the structure in (3.11) has an additional attachment point. A discourse grammar based on this type of elementary structure will predict more possible readings and might be useful in accounting for subtle differences between readings of a given discourse.

$\mathcal{A}16\ \forall k[imp(k) \leftrightarrow term(k) \wedge \neg lex(k)]$

We make sure that all terminal nodes are marked as such in the elementary tree descriptions in $\mathcal{W}$, and that lexemes occurring in sentences introduce lexical nodes in the input descriptions that represent them. As in nonincremental LDG, it is an integrity constraint of the grammar that lexical nodes are identified with the anchors in the lexical descriptions they select. Thus, all lexical nodes are terminals, but not every terminal needs to be lexical.

The *lexicalisation principle* now applies to terminal nodes, as it applied to lexical nodes before: all nodes in a verifying tree model of a discourse description must be positively and negatively anchored to a terminal node in the tree. At the same time, the discourse grammar shall be sufficiently relaxed to allow for implicit anchors occuring in between lexically consecutive constituents in the described discourse structure (see section 3.5.3 and 3.5.4).

Whenever constraints in LDG are relaxed for some reason, there is a danger that the resulting formalism is too tolerant, in other words, that it allows for unwanted verifying models or interpretations of discourse. Indeed, the occurrence of implicit anchors in the described discourse tree structure must be restricted to cases which can be linguistically motivated. Even though we want to allow for implicit anchors, we do not want implicit nodes and structure to appear just anywhere in the models of discourse descriptions. Further research is needed to find out if there are other linguistic constructs for which an analysis in terms of implicit anchors, or deduced nodes and structure, might be appropriate. For the time being we assume that implicit anchors occur only with discourse relations, that is, all implicit anchors are of word class 'rel':

$\mathcal{A}17\ \forall k[imp(k) \rightarrow rel(k)]$

3.5 Incremental LDG

While for sentence analysis both incremental and nonincremental models can be used, for a model of discourse processing only an incremental formalism comes into question. Discourses do not come as finished wholes. They are built up utterance by utterance, and each utterance is interpretable to the language user, given the discourse processed so far. The discourse model should reflect this. As we saw in the preceding chapter, however, the LDG proposed in Muskens (2001) is not incremental: input descriptions sum up all the lexemes that are relevant for the parsing of a sentence, and sentences are parsed as a whole.

As can be seen from Kallmeyer (1999), Egg et al. (1998), Gardent and Webber (1998), and Duchier and Gardent (1999) there are some incremental formulations of description grammar on the market. While in terms of expressive power the description grammar for discourse of Gardent and Webber (1998) would do equally

well, we prefer to provide an incremental variant of Muskens' LDG here. The benefit in this is that the incremental grammar formalism will inherit the fine-tuned anchoring system of nonincremental LDG, which makes parsing into a simple pairing off of positively and negatively marked nodes. Thus, parsing results directly in the set of linguistically relevant verifying tree structures. There is no need for an additional computation on the basis of which only the *minimal* models among those verifying the descriptions are selected, as is the case with the description grammars of Vijay-Shanker (1992), Kallmeyer (1999) and Duchier and Gardent (1999).

The structure of this section is as follows. In section 3.5.1 it is argued that the manner in which lexical nodes are circumscribed in Muskens' LDG leads to problems in an incremental grammar. We then propose an alternative way of circumscribing the lexical nodes in the models. This results in some adjustments in the form of the input descriptions (section 3.5.2), and in the formulation of what we shall call the left- and right-closure properties of discourse descriptions, (section 3.5.4). All this interacts in a certain way with the incrementation operation, which is defined in section 3.5.3.

As before, parsing is deduction, or rather, parsing is reasoning about discourse descriptions and the tree models that verify them. We go through an example of this in section 3.5.4. The example features an implicit discourse relation. Given his grammatical knowledge, the hearer is predicted to deduce that there is a discourse relation linking the tree representations of the two sentences. Furthermore, employing world knowledge and default reasoning, he is able to infer a specific, semantically appropriate discourse relation. Thus, the discourse is predicted to be structurally and semantically coherent. In two subsequent sections, some further properties of incrementation as it is modeled in our framework are discussed.

3.5.1 Circumscription of Lexical Nodes

As we just said, an attractive feature of the LDG formalism discussed in the preceding chapter is that the tree structures which verify a description are precisely those that represent an intuitively plausible linguistic analysis of a given sentence, i.e. they are the minimal models of the description. One factor among the properties of the grammar formalism which cause this effect is the circumscription of lexical nodes in the input descriptions and the circumscription of nonterminal nodes anchored to the lexical ones in the elementary descriptions. Circumscription ensures that there are no other nodes in the described tree structure or model than the ones that are explicitly named in the description.

It is not difficult to see that circumscription, as it is implemented in LDG, leads to problems in an incremental version of the theory. Consider the input description of the sentence 'John left early'.

(3.12) $\exists k_1 k_2 k_3 (k_1 \prec k_2 \prec k_3 \wedge \textsc{john}(k_1) \wedge \textsc{left}(k_2) \wedge \textsc{early}(k_3)$

$$\wedge \forall k(lex(k) \leftrightarrow (k = k_1 \vee k = k_2 \vee k = k_3)))$$

The circumscription clause $\forall k(lex(k) \leftrightarrow (k{=}k_1 \vee k{=}k_2 \vee k{=}k_3))$ says that there are no other lexical nodes in the tree described than the ones summed up in the input description. This is unproblematic as long as we consider just single input descriptions. In the discourse model, however, the grammar builds up a discourse description from a sequence of input descriptions, adding an input description per processed utterance. Obviously, with more than one input description, each one stating that there are no other lexical nodes than the ones it sums up, the combination of these descriptions is inconsistent given the axioms of the grammar. So having the circumscription clause in each of the input descriptions of the utterances in a discourse is too strong. Intuitively, we only want to say that there are no other nodes in a discourse tree than the ones that are named in the discourse description in total, or in the input descriptions gathered so far *collectively*. How to maintain the constraining effect of the circumscription of lexical nodes, while dropping the circumscription clause in each of the input descriptions?

The essential observation is that a discourse, or more precisely, an initial stretch of discourse, can be viewed as a string of contiguous lexical elements with a first and a last element.[12] It is perfectly reasonable to assume that after each incrementation step, there are no lexical elements in between the ones that were parsed so far, and that there is no lexical element which precedes the leftmost one. Furthermore, as the language user is able to obtain an interpretation of the discourse as if it were now finished, we take it that he may temporarily assume that there is no lexical element which succeeds the rightmost lexeme parsed so far. When another utterance follows, this assumption must of course be dropped, but after each incrementation step it can be added again.

What we have just formulated is a method to dynamically circumscribe the lexical nodes in the described discourse tree. This is an alternative to circumscription per input description which does not hamper the stepwise incrementation of discourse descriptions. In the following sections we will implement this in the grammar formalism. We shall adjust the input descriptions, define the incrementation operation, characterize the first and the last element in a discourse description, and go through an example of discourse parsing or reasoning about discourse descriptions given the resulting, incremental LDG.

3.5.2 Input Descriptions

To begin, we make sure that the lexical elements that are listed in input descriptions are immediate successors. For this, we introduce the relation of *immediate precedence*, denoted by the symbol $\prec_l$.

[12] This presupposes strict linearisation of the utterances in a discourse.

Immediate Precedence $\forall k_1, k_2[k_1 \prec_l k_2 \leftrightarrow k_1 \prec k_2 \wedge \neg \exists k_3(k_1 \prec k_3 \prec k_2)]$

Modifying an integrity constraint of the grammar on the form of input descriptions, we now assume that in any input description, the lexical nodes are ordered by $\prec_l$ rather than $\prec$. At the same time the circumscription clause is dropped. Since it was a side-effect of the circumscription clause to identify the lexical nodes as such, we add independent statements to that effect. So instead of (3.12) we now have

(3.13) $\exists k_1 k_2 k_3(k_1 \prec_l k_2 \prec_l k_3 \wedge lex(k_1) \wedge lex(k_2) \wedge lex(k_3) \wedge$
$\textsc{john}(k_1) \wedge \textsc{left}(k_2) \wedge \textsc{early}(k_3))$

This is weaker than the original circumscription clause, but it still excludes nodes in the model occurring in between the ones representing parsed lexical elements.

Furthermore, we assume that when a discourse description is incremented with an input description, witnesses are chosen for the variables in the input description. Moreover, for every input description or update *new* witnesses are chosen, that is, witnesses which have not been used earlier in the ongoing parsing process. Intuitively this is not an implausible move, since each input description represents a particular occurrence of a sentence uttered in the discourse under consideration. As in the preceding chapter we use numerals as constants of type ν (tree node) for witnesses.[13] For example, we would have

(3.14) $0 \prec_l 1 \prec_l 2 \wedge lex(0) \wedge lex(1) \wedge lex(2) \wedge \textsc{john}(0) \wedge \textsc{left}(1) \wedge \textsc{early}(2)$

The advantage of using witnesses is that at the time when incrementation takes place, input descriptions feature constants rather than existentially bound variables as node names. Consequently, the incrementation operation can refer to lexical nodes introduced in input descriptions that were processed earlier. This would not be possible if the corresponding node names were bound in their input description by an existential quantifier.[14]

[13]By convention we start counting at 0 and choose the next natural number $(i + 1)$ for the next node we encounter. The order of the numbers is not relevant, however. What matters is that for new nodes new witnesses are chosen, that is, numbers that are not in use yet.

[14]In this, and the following sections we provide a simple set-theoretic formulation of incrementation in LDG. The treatment is a bit sloppy, however, in that it involves some reasoning about witnesses at metalevel. This can be avoided by employing the following *dynamic alternative*. We distinguish between dynamic descriptions and static ones. **Dynamic descriptions** are of type $\nu(\nu t)$ while static ones are of type t. In input descriptions of the form $\lambda k_{in} \lambda k_{out}.\delta$, λk_{in} is the leftmost lexical node and λk_{out} the rightmost lexical node. For example, we have

(3.15) $\lambda k_1 \lambda k_3 \exists k_2(k_1 \prec_l k_2 \prec_l k_3 \wedge lex(k_1) \wedge lex(k_2) \wedge lex(k_3) \wedge$
$\textsc{john}(k_1) \wedge \textsc{left}(k_2) \wedge \textsc{early}(k_3))$

The initial discourse description is simply the first input description processed. Generally speaking, discourse descriptions are 'left-closed', that is, the discourse they describe starts with

3.5.3 The Incrementation Operation +

Discourse descriptions are sets of input descriptions, extended by means of an incrementation operation. The incrementation operation + does two things in particular. Firstly, it *adds* a new input description δ_{i+1} to the discourse description built up so far, Δ_i:

Addition $\Delta_i \cup \{\delta_{i+1}\} \subseteq \Delta_i + \delta_{i+1}$.

Secondly, for all utterances in a discourse except the first one, it makes the most recently processed utterance into the right neighbour or successor of the sequence of utterances that were processed before. More precisely, it makes the leftmost lexical node named in the new input description into the immediate lexical successor of the rightmost lexical node named in the discourse description. We define the relation of *immediate lexical succession (ILS)*, denoted by the symbol $\prec_\iota$ as follows.

ILS $\forall k_1, k_2[k_1 \prec_\iota k_2 \ \leftrightarrow \ (k_1 \prec_l k_2 \ \vee \ \exists k_3(k_1 \prec_l k_3 \prec_l k_2 \wedge imp(k_3)))]$

In other words, a node k_2 *immediately lexically succeeds* a node k_1, or k_2 is a *immediate lexical successor* of k_1 if and only if k_1 immediately precedes k_2, or there is a single implicit anchor which immediately succeeds k_1 and immediately precedes k_2. The effect is that when a pair of nodes is immediately lexically succeeding, the lexical yields of these nodes are consecutive. The *lexical yield* of a node k is the sequence of lexemes attached to the terminal nodes dominated by k. Only invisible implicit anchors can be in between. This is desirable with a view to the treatment of discourse relations that we envisage.

---

a particular leftmost lexical element, and there is a node in the described discourse tree which represents this leftmost lexeme throughout the processing of the discourse. We employ 0 as a name for the node representing this element in discourse trees. No verifying tree contains a node k which precedes 0. **Left-closure** consists in applying type $\nu(\nu t)$ discourse descriptions to 0: $\Delta(0)$. The **incrementation operation** + links dynamic input description to left-closed discourse descriptions. So for a left-closed dynamic discourse description Δ of type (νt) and a dynamic input description δ of type $\nu(\nu t)$ we have

$\Delta + \delta = \lambda k.\exists k_1\, k_2[k_1 \prec_\iota k_2 \wedge \Delta(k_1) \wedge \delta(k_2)(k)]$

The output discourse description is a conjunction of the input discourse description and the new input description. Furthermore, the rightmost lexical node named in the input discourse description is made into the immediate left neighbour of the leftmost lexical node named in the input sentence description. (See the main text for the precise definition of the $\prec_\iota$ relation.) Temporary **right-closure** of discourse descriptions consists in applying the discourse description to the node currently representing the rightmost lexeme in the discourse: $\Delta(n_{fin})$, where n_{fin} designates the node name or witness of the node representing the last lexeme. Left-closed discourse descriptions can also be turned static by existentially binding the output element: $\exists k \Delta(k)$. For static discourse descriptions, verifying tree modes can be obtained in the usual way.

As we said, the incrementation operation makes the leftmost lexical node named in the new input description into the immediate lexical successor of the rightmost lexical node named in the discourse description. We call this property of incrementation *Lexical Contiguity*. Using $n_{RML\Delta_i}$ to denote the witness of the node name in Δ_i representing the rightmost lexeme in the discourse, and $n_{LML\delta_{i+1}}$ to denote the witness of the node name in δ_{i+1} representing the leftmost lexeme of the input sentence, we define it as follows.

Lexical Contiguity $\{n_{RML\Delta_i} \prec_\iota n_{LML\delta_{i+1}}\} \subset \Delta_i + \delta_{i+1}, \quad$ where $i > 0$.

For obvious reasons Lexical Contiguity does not apply to the input description representing the first utterance of a discourse. That description updates the initial discourse description Δ_0, where Δ_0 is defined as the empty set:

Initial Discourse Description $\Delta_0 = \emptyset$

Interestingly, the initial discourse description is verified by an infinite number of different tree structures, namely all the tree structures allowed for by $\mathcal{G}$. So, in fact, the initial discourse description is a totally underspecified representation of any possible discourse which is grammatical according to $\mathcal{G}$. Every update of Δ_0 eliminates some of the possible tree structures or models allowed for by Δ_0 given $\mathcal{G}$. Subsequent updates eliminate further models.

Together, Addition and Lexical Contiguity characterize the incrementation operation +. Given the distinction between the first update and all subsequent ones, we define it as follows.

Incrementation

1. $\Delta_0 + \delta_1 = \Delta_0 \cup \{\delta_1\}$
2. $\Delta_i + \delta_{i+1} = \Delta_i \cup \{\delta_{i+1}\} \cup \{k_{RML\Delta_i} \prec_\iota k_{LML\delta_{i+1}}\}, \quad$ where $i > 0$.

To illustrate the above, consider the following discourse.

(3.16) *John left early.*
He was ill.

Processing the first utterance results in the following input description. Witnesses have been chosen for the node names.

(3.17) $0 \prec_l 1 \prec_l 2 \wedge lex(0) \wedge lex(1) \wedge lex(2) \wedge \textsc{john}(0) \wedge \textsc{left}(1) \wedge \textsc{early}(2)$

Updating the initial discourse description Δ_0 with (3.17) results in Δ_1:

(3.18) $\{0 \prec_l 1 \prec_l 2 \wedge lex(0) \wedge lex(1) \wedge lex(2) \wedge \textsc{john}(0) \wedge \textsc{left}(1) \wedge \textsc{early}(2)\}$

Processing the next utterance of the discourse results in the following input description. Notice that new witnesses have been chosen.[15]

(3.19) $9 \prec_l 10 \prec_l 11 \wedge lex(9) \wedge lex(10) \wedge lex(11) \wedge \textsc{he}(9) \wedge \textsc{was}(10) \wedge \textsc{ill}(11)$

Updating Δ_1 with (3.19) produces the discourse description Δ_2:

(3.20) $\{0 \prec_l 1 \prec_l 2 \wedge lex(0) \wedge lex(1) \wedge lex(2) \wedge \textsc{john}(0) \wedge \textsc{left}(1) \wedge \textsc{early}(2)\}$
$\cup \{ 9 \prec_l 10 \prec_l 11 \wedge lex(9) \wedge lex(10) \wedge lex(11) \wedge \textsc{he}(9) \wedge \textsc{was}(10) \wedge \textsc{ill}(11)\}$
$\cup \{2 \prec_\iota 9\}$

Since $2 \prec_\iota 9$ reduces to either $2 \prec_l 9$, or $2 \prec_l n \prec_l 9$ for some implicit node n, all terminal nodes named in Δ_2 are strictly ordered by immediate precedence. Verifying tree structures of Δ_2 cannot contain unwanted additional nodes in between the terminal ones named in the discourse description. This holds for all discourse descriptions.

Given Lexical Contiguity of input descriptions in general, it is ensured that if a newly processed sentence is not linked to its discourse context through an overt discourse operator incorporated in the sentence, it is related to the discourse context through (an implicit) discourse relation. There are no other options. Furthermore, it is a property of the grammar that discourse relations cannot be stacked, that is, if a sentence is linked to the discourse context through a discourse relation, there is only one of them. This will be illustrated in the next section. On the whole, every newly processed sentence is predicted to be structurally linked to its discourse context, either via an explicit discourse operator or via a discourse relation. We will refer to this property of incremental LDG as *structural coherence*.

3.5.4 Reasoning about Discourse Descriptions

So far, we have made sure that there are no lexical nodes intervening between the lexical nodes that are explicitly named in the discourse decription. However, as yet it is still possible that a verifying model contains lexical nodes preceding or succeeding the string of nodes representing the lexical elements actually occurring in the discourse. To prevent this, we introduce the following two axioms or general descriptions in $\mathcal{G}$. We assume that the language user is able to recognise the first and the last lexical element in a discourse, and we reserve the special node names 0 and n_{fin} for the representation of these elements. The axioms are called 'left-closure' (LC) and 'right-closure' (RC) respectively.

[15] In chosing the witnesses for the lexical nodes in (3.19) some numbers have been skipped. This is because, in fact, the missing numbers are employed as witnesses for the node names of nonterminal nodes named in the sentence description of the preceding sentence, see the more extended discussion of this example in the next section.

LC $\neg \exists k(k \prec 0)$

RC $\neg \exists k(n_{fin} \prec k)$

For any well-formed discourse description, axiom **LC** tells the language user that there is no verifying tree model which contains nodes preceding the node named 0. That node is the left-most lexical one in the described tree structure. Discourse descriptions which contain a leftmost lexical node named 0 will be called *left-closed*.[16] Similarly, axiom **RC** implies that there is no verifying model which contains nodes succeeding the node named n_{fin}, which is the rightmost lexical one in the described tree structure. Discourse representations containing a statement which identifies the rightmost lexical node named in the discourse description with n_{fin} will be called *right-closed*.

When a discourse description is both left- and right-closed, the lexical nodes in a verifying discourse tree must be precisely those specified in the description. Thus, with the help of Lexical Contiguity, left-closure and right-closure, the effect of the circumscription clause in nonincremental LDG is obtained: the lexical nodes in the verifying tree structures are circumscribed. Verifying trees of the discourse description can be obtained in the same manner as in nonincremental LDG, by pairing off positively and negatively marked nodes.

What is different now is that discourse descriptions are built up incrementally. Crucially, only discourse descriptions that are not right-closed can be incremented without becoming inconsistent. At the same time, after each incrementation step the language user is able to deduce the semantic content of the discourse processed so far, and reason about its possible disambiguations. In order to obtain the verifying tree models of the discourse processed so far, however, the update effect of possible continuations of the discourse must be ignored, that is, the discourse description must temporarily be right-closed.[17]

Thus we propose that after each incrementation step, the language user may assume that the discourse is finished and reason under the assumption that the discourse description is right-closed. Adding the hypothetical end statement $n_{RML\Delta} = n_{fin}$ to his premises, he can deduce node identifications and obtain the verifying trees of the current discourse description. When the discourse continues he drops the right-closure hypothesis. In general, when considering the interpretation of the discourse processed so far, we are interested in statements ϕ such

[16] In Van Leusen and Muskens (2003) a slightly different approach is followed. It is assumed that the root node in a discourse description is anchored to a special lexical element, the start element, and the start element is characterised by the left closure property. A pleasant consequence of this is that the root node comes with a lexical anchor, like all other nodes in a tree description. Furthermore, the first incrementation step need no longer be treated as an exception to Lexical Contiguity. Apart from this, the effect of the two approaches is exactly the same.

[17] Note that anything stated in the processed part of a discourse might be retracted later on through corrective discourse contributions, so nothing can be concluded from a discourse description if we take into account its possible continuations.

that $\mathcal{G}, \Delta \models (n_{RML\Delta} = n_{\mathit{fin}}) \rightarrow \phi$, where Δ is the current discourse description and $n_{RML\Delta}$ is the rightmost lexical node named in Δ.

To illustrate this, consider again the discourse description Δ_1, resulting from the language user's processing of the utterance 'John left early.'

(3.21) $\{0 \prec_l 1 \prec_l 2 \wedge lex(0) \wedge lex(1) \wedge lex(2) \wedge \textsc{john}(0) \wedge \textsc{left}(1) \wedge \textsc{early}(2)\}$

Given **LC** the language user knows that Δ_1 is left-closed. Using his lexical knowledge he will deduce the appropriate lexical descriptions for the lexemes in the input description.

(3.22) Δ_1: r^-

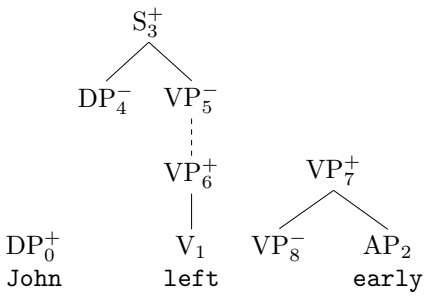

On the temporary assumption that the discourse is currently finished, the language user identifies the right-most lexical element named in the description with n_{fin}, adding the hypothetical end-statement

(3.23) $2 = n_{\mathit{fin}}$

to the discourse description. Δ_1 is now right-closed. We indicate this by means of a superscript ^{rcl} at the description label.

(3.24) Δ_1^{rcl}: r^-

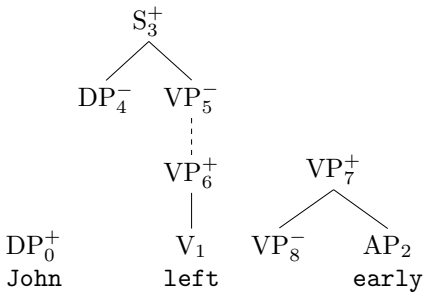

The hearer may now reason about the verifying models of the discourse description, mapping positively to negatively marked nodes. This results in

(3.25) Δ_1^{rcl}:

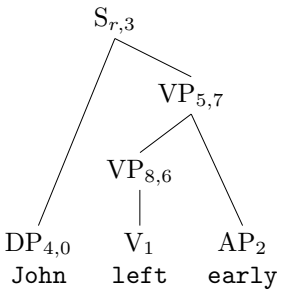

For Δ_1^{rcl} there is only verifying tree model; there is no ambiguity in the discourse parsed so far. When next 'He was ill.' is uttered, the hearer finds out that the discourse was not finished yet and drops the hypothetical end-statement. Δ_1 is incremented with the new input description and the appropriate lexical descriptions are selected, resulting in Δ_2.

(3.26) Δ_2:

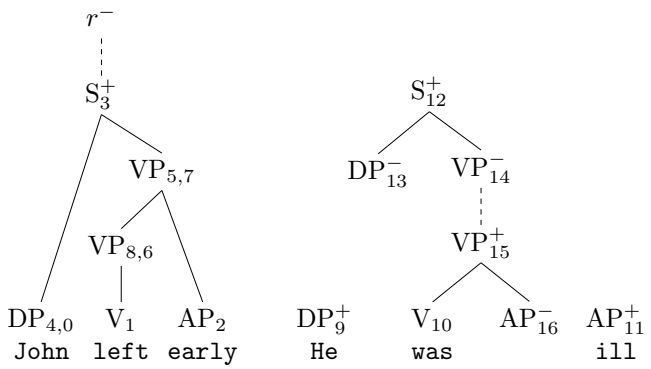

Notice that the completed parse of the sentence 'John left early.' as represented in (3.25) is transferred to (3.26), while this parse resulted from reasoning under the right-closure assumption. The question becomes important how much of a hearer's reasoning is independent from hypothetical end statements such as (3.23). This independent part can monotonically be transferred to the next phase of the reasoning process. The matter depends essentially on the basic structural properties of elementary tree descriptions, and the 'attachment potential' allowed them by the axioms or general descriptions in the grammar formalism. Though the issue is important it is currently still open.

As for conclusions which depend crucially on the right-closure assumption, we would suggest that the transfer of those can be thought of as the result of a bias or preference of the language user to preserve conclusions once drawn, or to minimize information loss. Unlike conclusions that are drawn independently from right-closure, they are cancellable.

Consider (3.26) again. The newly added input description now represents the last utterance in the discourse. The language user wants to parse it and find out what it means relative to its discourse context. The last word of the newly processed utterance is now the final word of the discourse, so he assumes

(3.27) $11 = n_{fin}$

and he reasons about the verifying tree models of the discourse description. First, reasoning about node identifications, he may parse the new sentence.

(3.28) Δ_2^{rcl}:

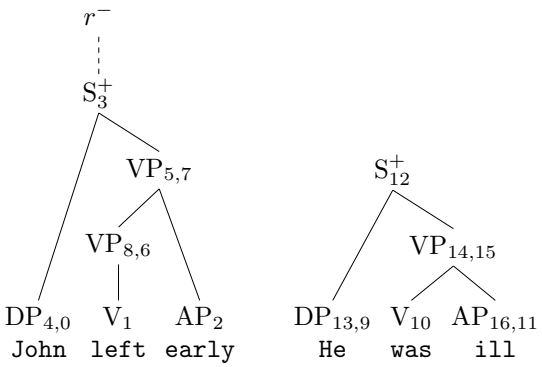

But then, what about the discourse as a whole? The description of the new sentence does not come with some overt discourse operator which introduces a negatively marked node to the left of the sentence. And no verifying trees can be obtained by identifying r with 3, or r with 12. In the absence of additional nodes and structure no verifying models for the discourse description can be found. There must be some—currently underspecified—additional nodes and structure, linking the nodes named 3 and 12.

In the reasoning process that leads up to this conclusion, there are two options to consider. Given Lexical Contiguity of the incrementation operation $2 \prec_\iota 9$, so either $2 \prec_l 9$, or $2 \prec_l n \prec_l 9$ and n is an implicit anchor. First, suppose $2 \prec_l 9$. Given left and right closure of the discourse description, the presence of any additional nodes in between the terminal ones named in the discourse description is excluded. As all terminal nodes in a tree structure are ordered by precedence, this excludes any additional terminal nodes. Given the anchoring axioms of the grammar, there can be no additional nonterminal nodes in the described tree structures either. But then, it is not possible to pair off positively and negatively marked nodes in Δ_2^{rcl} in such a way that all nodes become satisfied. So no verifying tree model is obtained.

The other option is that $2 \prec_l n \prec_l 9$ and n is an implicit anchor. According to axiom $\mathcal{A}17$, specified in section 3.4, implicit anchors occur only with discourse relations. Discourse relations, like lexemes, come with an elementary tree structure of the appropriate type. The lexicon $\mathcal{W}$ contains two elementary descriptions for elements of type discourse relation: one for coordinating and one for subordinating discourse relations. There are no other options. The language user may conclude that either the one or the other applies, and that the left and right argument of the discourse relation dominate the nodes representing the two sentences in the discourse, respectively. The following pictures show the two possibilities.

(3.29) Δ_2^{rcl} *(a):*

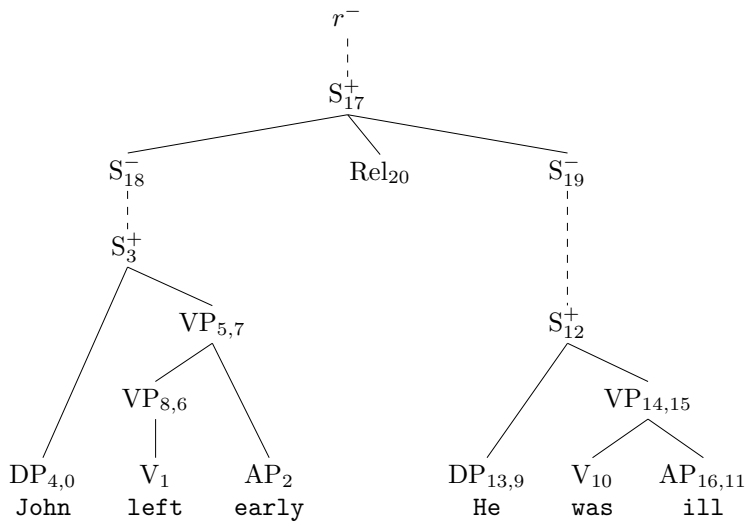

(3.30) Δ_2^{rcl} *(b):*

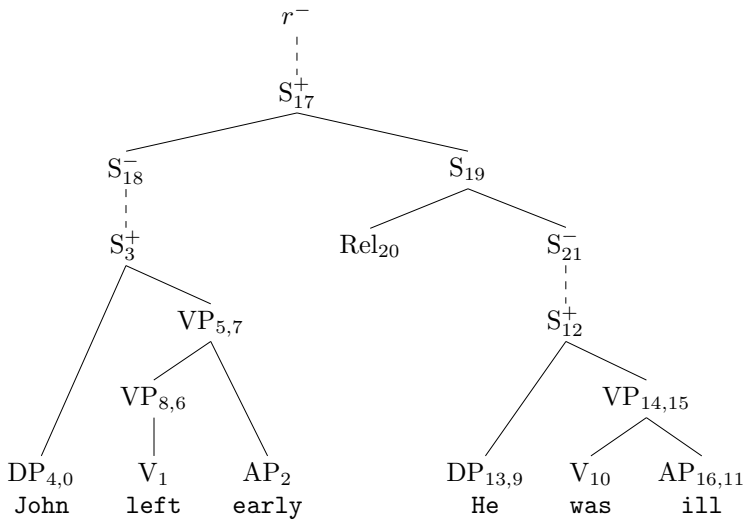

The inferred implicit anchor is the node named 20. Since in both Δ_2^{rcl}(a) and Δ_2^{rcl}(b), $2 \prec_l 20 \prec_l 9$, it can immediately be seen that no further implicit discourse relations can be deduced in between 2 and 9: the immediate precedence of 2 and 20, or 20 and 9, would be violated. For each of the possibilities in (3.29) and (3.30) a verifying model can be obtained by mapping positively to negatively marked nodes in the obvious way. Thus, the discourse is predicted to be structurally coherent.

At this point in the reasoning process, however, the discourse relation linking the second sentence to its discourse context is still both structurally and semantically underspecified. The grammar $\mathcal{G}$ enforces some measure of *semantic nontriviality* in the sense that, whichever the underspecified discourse relation is, it must be one of the finite collection described in the lexicon of the grammar. Employing his world knowledge $\mathcal{K}$ and default reasoning component ■, the hearer may determine a most preferred one among these. As in all other discourse theo-

ries, the selection of a semantically *most coherent* interpretation builds essentially on interaction with the hearer's preference system.

In the case of Δ_2^{rcl}, he may infer that the reason why John left early was that he was ill, that is 'He was ill' is linked to 'John left early' by a relation of explanation. There may be other possible readings, for example one in which John left early and he was ill, but there is no causal link. In that case, the two propositions might be related by a relation of *contination* or one of *background*. We assume that the hearer's default reasoning and preference system generates a preference for the explanatory reading. Because *explanation* is a subordinating discourse relation, the preference induces not just semantic, but also structural disambiguation. Option Δ_2^{rcl}(a) is now excluded, and we have

(3.31) Δ_2^{rcl}:

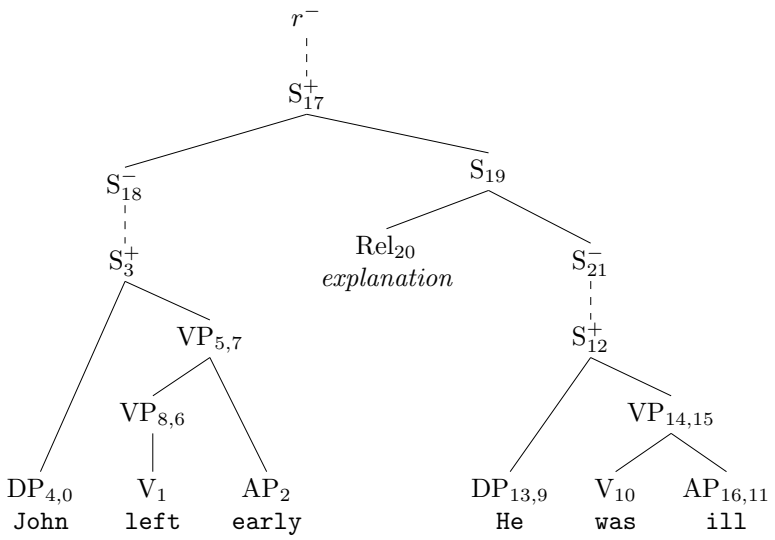

The default inference that the coherence relation is a relation of explanation is a prototypical example of a conclusion which will be transferred to the next phase of the reasoning process, but which might be cancelled when more information is available to the hearer, for instance through the analysis of subsequent discourse. Notice that if it would be cancelled, the two structural alternatives in (3.30) and (3.29) would be available again. The conclusion that there is some implicit discourse relation linking the sentence represented by the node named 12 to its discourse context would be preserved, however.

There is a single most preferred verifying tree structure for Δ_2^{rcl}, namely

(3.32)

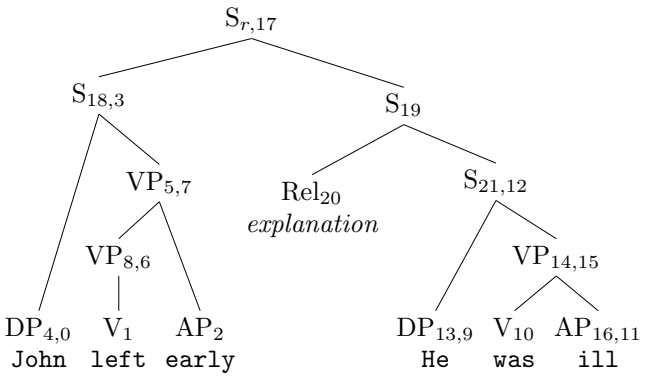

After each incrementation step, the language user may reason about the semantic contents of utterances and about the meaning of the discourse processed so far. As in the nonincremental LDG, each lexeme in the input comes with a semantic value and each nonterminal comes with a recipe for the composition of its semantics. Application of the recipe to the basic values produces the semantic value of the node under consideration. For any right-closed discourse description Δ_i^{rcl}, the discourse meaning will be defined as an update of the relevant context of interpretation with σ_r, the semantic value of the root. This will be explained in full in the next chapter, which zooms in on the semantic interpretation of discourse and the integration of felicity, or appropriateness conditions in the discourse theory.

3.5.5 Attachment at the Right Frontier

It is easy to see that the incrementation process influences the accessibility of the nodes named in discourse descriptions. We call a node named in a discourse description *structurally accessible* if it is a potential attachment site for a subsequent update of the discourse description.[18] A node named in a discourse description is a *potential attachment site* if it can be anchored consistently to a terminal node named in the input description resulting from a potential subsequent update. If the node is positively marked, it will find its negative anchor in the part of the tree structure representing subsequent discourse. If it is negatively marked, it will find its positive anchor in the part of the tree structure representing subsequent discourse. If it is unmarked, or 'satisfied' it cannot be an attachment site for subsequent updates.[19]

[18]In TAG for discourse and LDM structurally accessible nodes are often called 'open nodes' and structurally inaccessible ones 'closed nodes'. In description grammar approaches the term 'open node' is sometimes used to denote what we call negatively marked node names, and the term 'closed node' to denote both what we call positively marked node names' and 'satisfied node names'. We will avoid these terms or use them only in the first sense here, that is, 'open nodes' are structurally accessible ones and 'closed nodes' are structurally inaccessible ones.

[19]Though notice that unmarked nodes can be redescribed by subsequent updates in principle. There is no single constraint in the grammar which says that overlapping descriptions are prohibited. In practice, however, it is excluded through the interaction of the axioms of the

To illustrate, consider the following discourse description. It is the description which results when, after parsing and interpreting the discourse 'John left early. He was ill.', the hearer drops the right-closure assumption in expectation of the next discourse contribution.

(3.33) Δ_2:

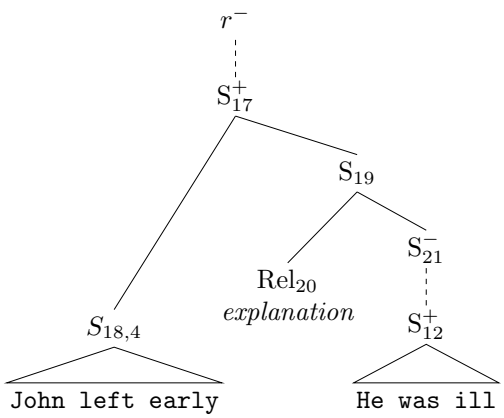

Given our definition of incrementation, subsequent updates place the tree described by the new input description immediately to the *right* of the structure already described by the input discourse description. Hence, the potential attachment sites of the discourse description in (3.33) are the nodes named r, 17, 21 and 12. All other nodes named in Δ_2 are structurally inaccessible or *closed off* for attachments of subsequent updates.

For example, in the updated discourse description represented in (3.34) below, it is impossible to identify the node named 18 with the node named 23 in the input description resulting from the update without violating the axioms of the grammar or the properties of the discourse description. Suppose $18 = 23$. Since nodes in a tree can only have a single mother node, it would follow that $17 = 22$, $19 = 24$, $20 = 25$, and $21 = 26$. However, 20 cannot be identified with 25 because the first strictly precedes the second. Hence, 18 cannot be identified with 23; it is structurally inaccessible to the update.

(3.34) Δ_3:

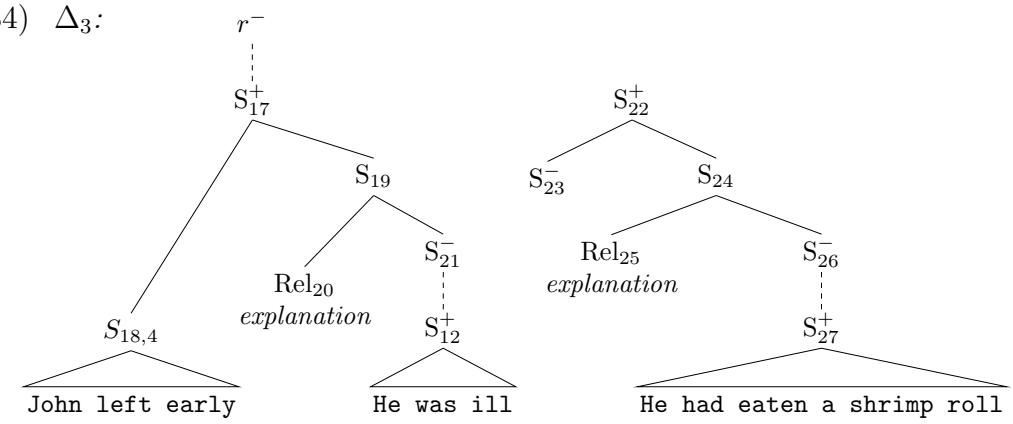

grammar and the properties of input descriptions of the incrementation operation.

For the new input description represented in (3.34) there are two possible attachments: either ($r \lhd^* 22$ and) $23 \lhd^* 17$, or $21 \lhd^* 22$ and $23 \lhd^* 12$. Under the assumption that the discourse description is right-closed, this comes down to $r = 22$ and $23 = 17$, or $21 = 22$ and $23 = 12$. The two attachments result in discourse descriptions representing different interpretations of the discourse. There is a form of ambiguity here, similar to the sentence internal structural ambiguity discussed in section 2.4.2 of the preceding chapter. Following Gardent and Webber (1998) and Duchier and Gardent (2001) we shall refer to this as *attachment ambiguity*: in general an update can attach at different sites in a described discourse tree. Different attachments come with verifying models that are structurally different, and result in different readings of the utterance that triggered the update. Additional semantic constraints in the grammar and preferences generated by the language user's default inferences cut down ambiguity.

For the purpose of the discussion, we have simplified the example a bit by already selecting the relation of *explanation* as the preferred implicit relation linking S_{27} to the discourse context. Furthermore, we assume the hearer prefers a reading in which the fact that John ate a shrimp roll explains that he was ill over one in which John's having eaten a shrimp roll explains that he left early, because he was ill. This corresponds to the attachment $21 \lhd^* 22$ and $23 \lhd^* 12$ in Δ_3:

(3.35) Δ_3:

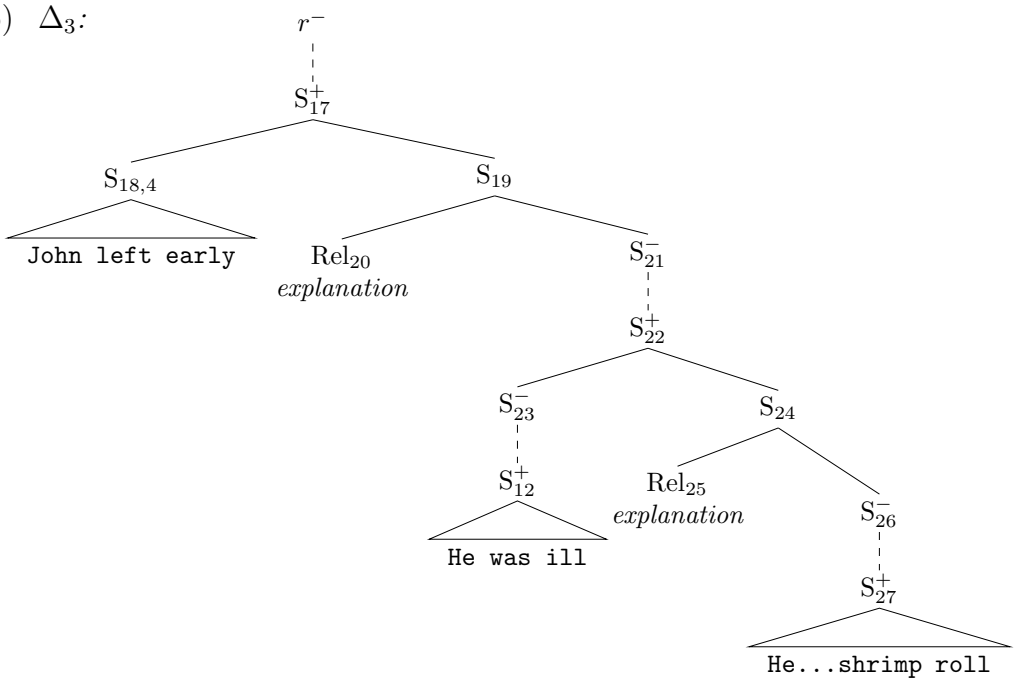

Independently of right-closure of the discourse description, the attachment of the update leads to the closing off of part of the discourse description. Whatever the attachment position of the next update is going to be, all nodes dominated by or identical to the node named 23 in (3.35) are now structurally inaccessible to

subsequent updates. In other words, the hearer may deduce

(3.36) Δ_3:

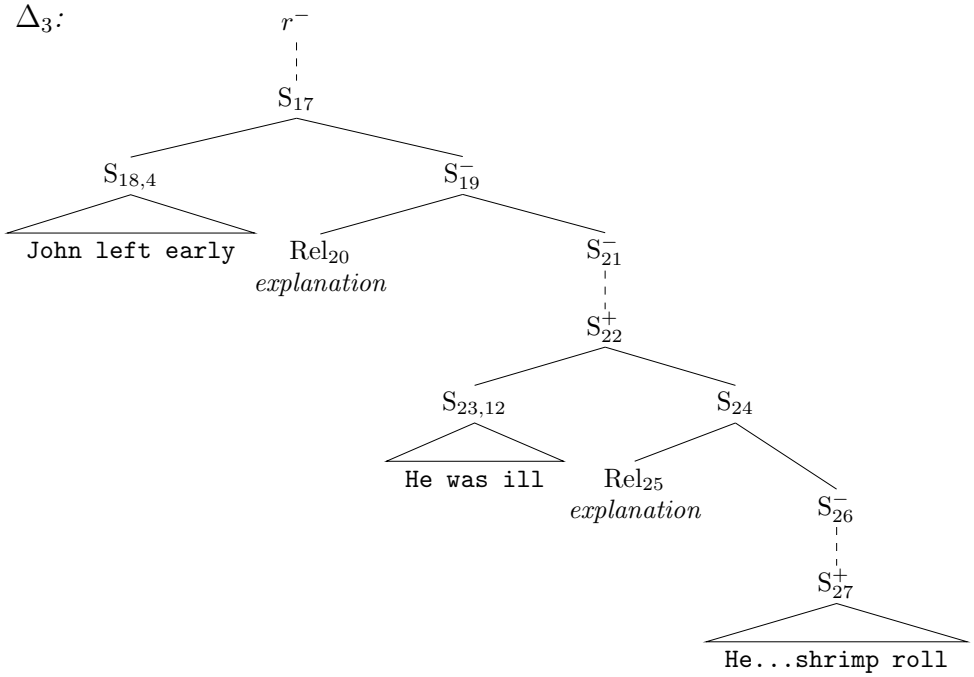

The potential attachment sites of this discourse description are the nodes named r, 17, 21, 22, 26 and 27. As the reader can see from the above series of pictures, there is something like a right frontier in LDG. In TAG for discourse and LDM the right frontier is roughly the right edge of the discourse parse tree, at which the adjunction of trees representing newly processed input takes place, see section 3.4. What is updated in LDG is not a discourse tree, however, but a discourse description, a collection of logical statements whose verifying models are tree structures. *Right frontier*, in LDG, is essentially a pictorial term for the set of potential attachment sites or structurally accessible nodes named in a discourse description. A given discourse description may be verified by more than one tree structure, and each of these tree structures comes with its own right frontier, in the sense of TAG for discourse or LDM. The discourse description underspecifies these right frontiers.

As in TAG for discourse and LDM, attachment at the right frontier is a consequence of the incrementality of the grammar. Due to the fact that, first, all input descriptions are anchored to at least one terminal element, second, all terminal elements are strictly ordered by precedence, and third, incrementation takes place from left to right, the nodes that are structurally accessible to a new update can only be nodes in what might be dubbed the 'right frontier zone' of the pictorial representation of the discourse description.

Exactly how far this zone stretches depends on the axioms of the grammar, and the integrity constraints on the form of elementary trees described in the lexicon. While in simple examples like the one discussed above, it is easy to show

which nodes in a given discourse description are structurally accessible it is a matter for further research what can be said about the attachment potential of discourse descriptions in general, given the formal properties of LDG. This relates directly to the issue of monotonic transfer of conclusions from one phase in the parsing process to the next, touched upon in section 3.5.4. It all concerns the inherent closure properties of discourse descriptions.

3.5.6 Deterministic Discourse Processing

The processing of an arbitrary discourse contribution results in an incrementation of the input discourse description. The incrementation operation +, defined in section 3.5.3, is monotonic: one or more logical statements are added to the set of statements that make up the input discourse description. Apart from the temporary right-closure statements, and, sometimes, conclusions drawn earlier on the basis of default reasoning, no statement is ever removed from a discourse description. Disregarding the effect of right closure and default reasoning, every update strengthens the discourse description and eliminates potential verifying models of the discourse. As we shall see in the next chapter, this goes for all updates, including those induced by corrective discourse contributions. While in terms of their semantic interpretation they are nonmonotic, they induce monotonic updates of discourse descriptions.

An incrementation step results in an update of the discourse description. The language user may subsequently reason about the *attachment* of the newly added input description. Once again, this is a monotonic operation: node names, dominance relations, and possibly node identifications are added to the discourse description, but no statements are retracted from it. Other than in TAG, where the adjunction operation changes properties of the parse tree it operates on, attachment in description grammar only adds information to the discourse description.

By virtue of the monotonicity of both the incrementation operation and the attachment of newly processed input descriptions, the discourse model allows for deterministic discourse processing in standard cases of attachment ambiguity, cooccurring with what might be called *argument underspecification*. In general, it is possible that the argument of a discourse relation is not completely specified yet at a given moment in processing the discourse, and that it is extended by a subsequent update. Consider

(3.37) a. *I'm not going to ask Jacky to take care of the children,*
because she always calls off at the last moment.

The current right argument of the connective 'because' is the proposition 'she always calls off at the last moment'. As the continuation in (3.38) shows, this is not necessarily the *final* right argument, or the whole reason why the speaker is not going to ask Jacky to take care of the children.

(3.38) a. *I'm not going to ask Jacky to take care of the children, because she always calls off at the last moment.*

b. *She isn't very nice to them either.*

The reason for not asking Jacky to take care of the children is not just that she always calls off at the last moment, but also that she is not very nice to the children. As is observed in Gardent and Webber (1998) and Duchier and Gardent (1999, 2001), argument underspecification is closely related to incrementality. Since the information conveyed in a discourse is only made available step by step, and arguments can be complex discourse units, it is unavoidable that an argument is not always specified completely in a single update. Argument underspecification is modeled directly in description grammar, through the underspecification of the dominance relation between 'current', and 'final' arguments of discourse relations. The following discourse description, resulting from processing (3.38a) illustrates this.

(3.39) Δ_1:

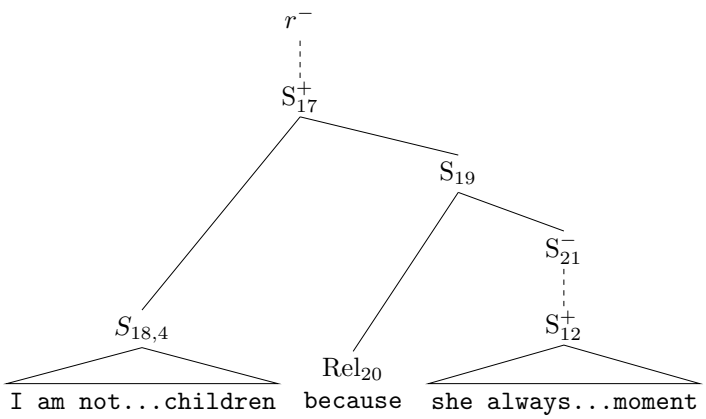

The node named 12 represents the current right argument of the causal connective, the node named 21 its final right argument. Reasoning under the right-closure assumption, the hearer will identify the two, and obtain the meaning of the discourse as it is after processing (3.38a) is obtained. When the discourse is continued, however, the current argument can be extended.

In Gardent and Webber (1998) and Duchier and Gardent (1999, 2001) it is pointed out that the description grammar treatment of argument underspecification and attachment ambiguity conforms to the *deterministic hypothesis* of Marcus et al. (1983). The hypothesis says that cases of attachment ambiguity that require no conscious reanalysis are processed deterministically. Indeed, the extension of an 'uncompleted' argument, as occurring in (3.38b), does not seem to lead to an increase in processing load. This seems to hold for cases of attachment ambiguity as well. Intuitively, no reinterpretation or backtracking takes place in the incrementation step. If this intuition is correct, then description grammar models incrementation in a cognitively plausible way.

Interestingly, there are also cases of ambiguity—so-called *garden-path ambiguity*—which do require conscious reanalysis. The deterministic hypothesis says that this type of ambiguity is not processed deterministically. Consider, for instance, the garden-path expression 'The horse raced ...' as occurring in the following two sentences.

(3.40) a. *The horse raced past the barn.*

b. *The horse raced past the barn fell.*

Sentence (3.40a) tells us that there is a horse which raced past the barn. Sentence (3.40b) says that there is horse that was raced past the barn, and which fell. In (3.40a) 'raced' is the finite verb, taking 'the horse' as its subject, in (3.40b) it is a participle heading the modifying phrase 'raced past the barn'. Especially in (3.40b), the hearer may get on the wrong track. Initially he may interpret 'raced' as the finite verb, and repair this only when at the end of the sentence 'fell' is encountered.

As is pointed out in Gardent and Webber (1998), garden-path ambiguity cannot be captured just by dominance underspecification in description grammar. It can, however, be represented by a disjunction over syntactic or semantic alternatives within the sentence description—this is a logical possibility in description grammar. And the ambiguity can be resolved by monotonically adding further information to the description, just like ambiguities represented by dominance underspecification. Thus, relative to these examples the grammar is not constraining enough: it does not predict that there is backtracking or indeterminacy in the parsing process.

As is suggested in Gardent and Webber (1998), if we want to account for examples of this kind, which lead to cognitive processing difficulties, we need an approach which combines underspecification with a preference system that highlights a specific reading: the reading which corresponds to the hearer's currently preferred interpretation. In other words, on the basis of such a preference system, or a set of 'biases', one of the disjunctive alternatives is chosen nondeterministically. When in the course of the parsing process the chosen option turns out to be incorrect, the preference is canceled and the previously dispreferred alternative pursued.

The treatment of garden-path sentences is one of the aspects of discourse interpretation that motivated us to make room for a preference system in the framework of discourse interpretation (see section 3.2). We assume that the language user generates preferences over alternative verifying models of discourse descriptions, on the basis of his world knowledge and default rules in ■.

As we have seen in the preceding sections, the preference system is applicable in a much wider domain than just garden-path sentences. Many cases of attachment ambiguity involve the choice of an appropriate, but underspecified discourse relation, linking the new input description to the discourse description. Some further preferences relevant for discourse processing are mentioned in Gardent and

Webber (1998) and Duchier and Gardent (2001). A well-known constraint in cognitive linguistic approaches is a bias towards 'low attachment': a preference for structures in which the incoming unit attaches 'low in the tree' and can be obtained by minimising the most recent dominance link.[20] It is certainly worthwile to investigate the implementation of this, and similar cognitively motivated constraints on discourse interpretation in LDG.

3.6 Summary and Loose Ends

Summarizing, in this chapter we put forward a framework of discourse interpretation which is built up around an incremental LDG formalism. The framework consists of an incremental grammar system G^+, a reservoir of world knowledge $\mathcal{K}$, and a default inference component ■. G^+ includes a discourse grammar $\mathcal{G}$. The discourse grammar represents the language user's grammatical knowledge, the reservoir of world knowledge models his nonlinguistic knowledge, and the default inference component models part of his reasoning capacity. G^+ generates logical descriptions of discourse. Discourse is processed stepwise, and each processing step induces an incrementation of the discourse description built up so far. Parsing is reasoning about discourse descriptions: the hearer reasons about the tree models that verify a discourse description given $\mathcal{G}$ and $\mathcal{K}$.

Generally speaking, discourse cannot be fully disambiguated on the basis of grammatical knowledge alone. Underspecification is ubiquitous at discourse level. Anaphoric elements must find some suitable antecedent, discourse connectives must find an appropriate attachment position, and discourse relations must be semantically specified. While the syntactic and semantic constraints in $\mathcal{G}$ determine the potential readings of a discourse, the hearer employs his knowledge of the world and default reasoning to make choices between these readings. We assume the hearer is able to weigh alternative readings against each other and opt for one or more preferred ones. The role of the default inference component ■ in the framework of discourse interpretation is to generate preferences over verifying models of discourse descriptions. In the presence of new information preferences may be dropped.

Thus, the framework implies a division of work in line with the architecture of discourse interpretation systems as proposed in both Webber et al. (1999) and Asher and Lascarides (2003). The grammar accounts for the lexical and compositional semantics of discourse. The semantics can be underspecified, however. Defeasible reasoning accounts for further disambiguation. Asher and Lascarides explicitly identify the interaction between compositional semantics and defeasible reasoning with the semantics-pragmatics interface: the task of *pragmatics* is to replace the underspecified elements with more complete information (p.113).

[20] See for example Crain and Steedman (1985), Frazier (1995), and Chen and Vijay-Shankar (1995).

Though this is obviously a theory-dependent characterisation of pragmatics, it is one sensible way of coming to grips with the notion. We will follow their suggestion and will henceforth use the term pragmatics in our framework for matters of discourse analysis that must be accounted for in terms of the language user's preferences and nonlinguistic knowledge.

The grammar system G^+ is an incremental variant of LDG as proposed in Muskens (2001). Discourse is processed stepwise, and each processing step induces an incrementation of the discourse description built up so far. The incrementation operation adds a new input description to the ones gathered so far and makes the tree described by the input description into the immediate right neighbour of the described discourse tree. The grammar predicts that any grammatical discourse is structurally coherent in the sense that every newly processed sentence is structurally linked to its discourse context, either through an overt discourse operator, or through an implicit discourse relation. Attachment takes place at the right frontier of the discourse description, that is, at some structurally accessible node named in the description. The selection of a semantically most coherent reading of a given discourse builds essentially on preferences.

The move to an incremental grammar system necessitated some changes in the LDG formalism. Instead of circumscription of lexical nodes per input description, the nodes representing lexical elements in the input descriptions are stated to be immediately consecutive. Furthermore, all discourse descriptions are left-closed, that is, there is a leftmost lexical node named in a discourse description, and this node cannot be preceded by any other node in verifying trees. Finally, after each incrementation step the language user may reason under the hypothesis that the last lexical element named in a discourse description is the rightmost lexical node in verifying trees of the discourse description. This was called hypothetical right-closure. When a discourse description is both left- and right-closed, the lexical nodes in a verifying model must be exactly those specified in the discourse description. As in nonincremental LDG, positively and negatively marked nodes can then be clicked together, and the verifying tree models of the discourse description are obtained.

The dynamic closure procedure raises an interesting question for further research: how much of the conclusions drawn on the basis of right-closed discourse descriptions can be moved monotonically from one discourse description to the next, after an incrementation step? To put it differently, what can be concluded from discourse descriptions that are left-closed but not right-closed?

Clearly it is desirable that certain properties should hold. Fortunately, this is the case for the structural coherence property, which plays a central role in the inference of implicit discourse relations. The conclusion that there is an implicit relation (if there is no overt discourse operator) is preserved under incrementation. The specification or choice of a particular discourse relation in a given discourse

is not necessarily preserved, however. This is the type of conclusion that builds on preferences or default reasoning; it can be cancelled when more information is available.

Secondly, we want the completed parse of clauses below sentence level and possibly also of complex sentences to be preserved under incrementation. Relative to this matter a further articulation of sentence theory will be relevant. Additional constraints defining clause bounderies, or rather clauses as maximal domains of interpretation for the quantifiers they contain, are necessary anyway.[21] Furthermore, it is sometimes claimed that sentencehood has a particular semantic impact of its own, for example on the interpretation of kataphors, so there may be reason to implement a notion of sentence-boundedness as well. In both cases the characterisation of a particular type of constituent as a local maximal domain will reduce the attachment possibilities of subsequent updates.

Apart from these considerations, it seems that the completion of discourse units is sometimes linguistically marked. The clearest example of this is with discourse units linked by the parallel (or list) relation, where intonational marking of the nonfinal elements of the list signals the continuation of the listing, while the final element in the parallel relation is intonationally marked as the last item. Furthermore, the end of a sentence is usually intonationally marked as such.[22] In each case, the intonational marking could be viewed as a trigger for the hypothetical closure of the part of the discourse description that covers the discourse unit under consideration.

Incrementation is a monotonic operation in the LDG system. Whatever the semantic update effect of a newly processed discourse contribution is, the incrementation of the discourse description it induces is a pure expansion: apart from hypothetical end statements and conclusions drawn on the basis of preferences, no information is retracted from a discourse description. Attachment of trees described by new input descriptions at what is commonly referred to as the right frontier of the discourse description is monotonic as well: node name identifications and dominance relations are added to the discourse description. As such, the treatment of incrementality in LDG conforms to the determistic hypothesis of Marcus et al. (1983): standard cases of attachment ambiguity, as induced by incremental discourse processing, do not require conscious reanalysis. They are processed deterministically in LDG.

The deterministic hypothesis also says that when conscious reanalysis or 'backtracking' does occur, discourse is not processed monotonically. This prediction applies to the treatment of garden-path sentences in LDG. As it turned out, however, the grammar system on its own is not sufficiently constraining to explain

[21] Compare Cimiano and Reyle (2005) on clause boundedness of quantifier scope in LDG/UDRT.

[22] See Pierrehumbert and Hirschberg (1990), on the form and interpretation of boundary tones.

the backtracking. Only in combination with default reasoning, which triggers a preference for a reading which later turns out to be unacceptable, the prediction is born out. Thus, the treatment of garden-path sentences provides some independent support for the idea that in one form or another a preference system or set of 'weak constraints' must be part of a framework of discourse interpretation.

The precise implementation of this awaits further research. What we are interested in in this dissertation is the function specification of the preference system: what are the kind of linguistic phenomena that could or should be explained in terms of biases or preferences? The distinctive feature, of course, is cancelability: preferences can be violated, and when this happens it does not lead to the prediction that the described discourse is ungrammatical. This sets them apart from the constraints that make up the grammar $\mathcal{G}$. These must be satisfied for the discourse to be grammatical. The grammar system processes discourse deterministically. Nondeterministic aspects of the interpretation process are to be accounted for in terms preferences or default conclusions.

Chapter 4

Discourse Semantics

4.1 Introduction

In the preceding chapter a framework for discourse interpretation was introduced. The focus was on the general properties of the framework, on discourse syntax and the characterisation of incremental Logical Description Grammar. In this chapter we unfold the semantic theory of the framework.

A central issue in discourse semantics is context dependence of interpretation. Discourse contributions are not made in an information vacuum. They relate in many ways to what was said and implicitly assumed in the preceding discourse, to the actual situation in which the discourse takes place, and to the world knowledge of the participants. The semantics of our framework must be arranged so that in principle such dependencies can be accounted for.

In answer to this demand we adopt a dynamic perspective on meaning, in which the meaning of a discourse contribution is viewed as the change it brings about in the context of interpretation (Kamp 1981; Heim 1982). We shall zoom in on the notion of *discourse meaning* in section 4.2, defining it as the update of what is implicitly presupposed at the start of the discourse with what is said. The discourse meaning resulting from an initial part of discourse serves as the context of interpretation for the next contribution in the discourse. Different classes of contributions may be distinguished according to their update effect and the requirements they put on the context of interpretation. Notably, we shall distinguish nonmonotonic discourse contributions, which result in revisions of the context of interpretation, from monotonic ones, which only expand it (section 4.3). General felicity conditions on consistent and informative updating can be specified in the semantic theory as general constraints on discourse meanings.

In the main body of this chapter (section 4.4) we will formulate these general constraints, but we will mainly be concerned with aspects of context dependence of interpretation that result from the use of particular elements or linguistic con-

structs in the language. Broadly speaking, such aspects of context dependence fall under the heading of anaphoricity, presupposition, information structure,[1] and coherence relations. Information structure will not be discussed in this work, but we will provide a treatment of various anaphors, presupposition triggers, discourse connectives and discourse relations. As in chapter 2, discourse relations are implicit elements in the lexicon of the grammar, which, like explicit connectives, describe or model relations of coherence in discourse. It shall be argued that all of the elements mentioned generate constraints on what we shall call, in line with Karttunen (1974), their *local context*. The local context of a linguistic expression is the semantic, contextual information that is accessible to it given its position in the discourse structure. We will define a notion of local context for each node in a discourse tree, and incorporate constraints on local contexts in the elementary descriptions of context sensitive elements of the language.

Thus, in our framework context sensitive constraints shall figure as an additional parameter in the compositional semantics of discourse. By incorporating conditions on local contexts in the lexical descriptions of context sensitive elements a straightforward declarative treatment of anaphora resolution, presupposition satisfaction, and appropriateness conditions of discourse operators will be obtained. In addition, through the simultaneous description of syntax and semantics in LDG there will be a smooth integration of grammatical levels.

A prerequisite for the account of discourse meaning and context sensitivity that we have in mind is that a less sketchy specification of the semantics featuring in discourse descriptions is provided than was given in the preceding chapters. Remember that the description language of the grammar framework is classical type logic $\mathcal{L}$. In this chapter we shall employ a sublanguage of $\mathcal{L}$ as semantic representation language within the linguistic descriptions. This will be a compositional DRT along the lines of Muskens (1996). In fact it will be a variant of compositional DRT that offers a notion of DRSs which is somewhat more finegrained than the one developed there. The finegrained DRT sublanguage of $\mathcal{L}$ is defined in appendix A.

The choice of compositional DRT for semantic representation language supports the comparability of our framework with other discourse theories employing DRT or underspecified DRT, notably SDRT (Asher and Lascarides 2003), and furthers the synthesis of linguistic analyses across these frameworks. The logical description method, however, is independent of the particular semantic formalism it is applied to here and other choices could be made.[2] Relatedly, our aim will be

[1] We use the term 'information structure' to refer to the subdivision of linguistic expressions (and their meanings) in parts with an informationaly different function, such as theme-rheme, focus-background, and topic-comment.

[2] Various alternative semantic representation languages figure in related linguistic models. For example, 'flat semantics' is used in the 'TAG for discourse' paradigm (Stone and Webber 1998); Vijay-Shanker's description grammar features logical forms (Vijay-Shanker 1992); in the Saarbrücken approach towards description grammar for discourse 'constraint language

not so much to develop 'new' theories of anaphora and presupposition resolution, or of discourse relations and coherence, but rather to show that such theories can be integrated in the framework of discourse interpretation simply through adding constraints to the grammar. As before, a language user will obtain a (possibly underspecified) interpretation of the discourse by reasoning about the verifying models of the discourse description given his grammar and his world knowledge. Thus, we hope to illustrate the functionality of the LDG framework as a logical toolkit for the specification of discourse theory in general.

We leave to further investigation the issue how participants change and extend their knowledge, attitudes, and commitments through conversation, and how this can best be treated in the LDG framework of discourse interpretation. The discourse semantics as it will be implemented will be defined in terms of the bare semantic contents of the participants' contributions. Discourse meanings are 'descriptions of reality' resulting from the sequence of contributions made in the discourse; they change with the participants' point of view and always reflect the most recent speaker's view of the reality described. They do not represent directly the commitments and attitudes of all participants at every moment in the discourse. We expect, however, that the framework can be extended to account for participants' commitments, either by adding inference rules so that participants' commitments can be inferred from the update history of the discourse, or by integrating commitments directly in semantic values and discourse meanings.

We conclude in section 4.5.

4.2 Discourse Meaning

We presuppose a framework of discourse interpretation $\mathcal{F}_a = \langle G^+, \mathcal{K}, \blacksquare \rangle$ of the kind developed in the preceding chapter. $\mathcal{F}_a$ models the knowledge of an interpreting agent or language user a, who can, but need not be a participant in the discourse under consideration. The incremental grammar system G^+ includes a description grammar for discourse $\mathcal{G}$, which represents the language user's grammatical knowledge of a given natural language. $\mathcal{K}$ represents his nonlinguistic, general knowledge or world knowledge, and $\blacksquare$ stands for his default reasoning capacity and capacity to assign preferences over verifying tree models of discourse descriptions. $\blacksquare$ can access both linguistic and nonlinguistic knowledge.

For any given stretch of discourse T the grammar system generates a discourse description Δ_T. This is a set of statements in a logical language describing the syntactic and semantic relations characterising the discourse processed so far. Discourse descriptions figure as underspecified representations of the discourse or *discourse contexts*. They represent the language user's knowledge of the discourse

for lambda structures' is used (Egg et al. 2001); the Linguistic Discourse Model employs a Montague-style semantics with unification variables (Prüst et al. 1994). Finally, SDRT (Asher and Lascarides 2003) employs segmented DRSs, a combination of DRT and rhetorical relations.

processed so far. In this chapter we shall introduce the notion of a local context. *Local contexts* are semantic objects attached to the nodes in a discourse tree. They represent, so to speak, the semantic dimension of the discourse context for each individual discourse unit.

On the basis of a left- and right-closed discourse description, the language user can reason about the tree models or linguistic structures that verify a discourse description given $\mathcal{G}$. In particular, he may reason about the semantic content of the sentences uttered, and the meaning of the discourse processed so far. The (possibly underspecified) *semantic content* of the discourse is represented by σ_r, the semantic value of the root node in a discourse description. This value is built up compositionally from the semantic values of the clauses or sentences uttered in the discourse. The *meaning* of a discourse will be defined as $B \oplus \sigma_r$, the update of an 'implicit background' or main local context B with σ_r.[3]

From a dynamic point of view, the *implicit background* can be thought of as the initial context of interpretation. From a static point of view it should be thought of as consisting of whatever is taken for granted or presupposed[4] by the participants during the conversation in question, according to the interpreting agent a. The background is highly underspecified, as what is presupposed is open for negotiation between discussion participants. Hence, we speak of the *implicit* background. In the course of the conversation it is specified, or as it were, constructed, through the linguistic presuppositions or context sensitive constraints evoked by the discourse.

The implicit background must be distinguished from the interpreting agent's reservoir of world knowledge $\mathcal{K}$. B is a semantic object in a discourse description, $\mathcal{K}$ is not. While B models what counts as presupposed *given the current discourse* in the eyes of the interpreting agent a, $\mathcal{K}$ represents his private world knowledge. B is not necessarily completely supported by all of the participants at every moment in the discourse, given their private knowledge and beliefs. In contrast, $\mathcal{K}$ is always supported by the interpreting agent a. When a is a participant in the conversation there is a systematic interaction between the linguistic presuppositions attached to his contributions and $\mathcal{K}$. For example, if a is truthful, he will not confirm a statement which carries presuppositional information which is in conflict with $\mathcal{K}$. Rather, the conflict will motivate him to utter a correction or to query the statement made. Rules describing such interactions may be specified within the discourse model, but this will not be done in this dissertation.

A language user a may *evaluate* a discourse contribution resulting in a discourse description Δ by hypothetically updating his world knowledge $\mathcal{K}$ with $(B \oplus \sigma_r)^{\Delta}$, the possibly underspecified discourse meaning determined by Δ. On the basis of the hypothetical update he may generate preferences over the verify-

[3] $\oplus$ is a merge operator conjugating DRS's, see section 4.4.1.

[4] This is *pragmatic* presupposition or 'common background belief', in the sense of Stalnaker (1974).

ing trees of Δ, employing his default reasoning engine $\blacksquare$.

Participants' attitudes and commitments In our framework of discourse interpretation we concentrate on semantic contents of utterances rather than the participants' attitudes towards that content. Discourse descriptions describe how semantic contents relate to each other, and discourse meanings are determined by them. As participants may disagree on certain points, it is not necessarily the case that all of the semantic content of a discourse is believed or supported by each of the participants at any point in the conversation. Rather, the semantic content of a discourse represents what is currently proposed for acceptance to the hearer by the most recent speaker. Our discourse model does not tell directly how a given discourse contribution changes the attitudes or commitments[5] of the participants in a given conversation, but it defines, somewhat more abstractly, pre- and postconditions of acceptance on discourse contributions made relative to arbitrary contexts of interpretation.

We embrace a *dynamic* notion of discourse meaning.[6] Discourse contributions in our model result in 'snapshots' or descriptions of reality representing the most recent speaker's point of view. With discourse contributions such as denials and corrections, the discourse meaning changes nonmonotonically: the previous snapshot is lost and cannot be retrieved from the new one. To illustrate, consider the following discourse, in which participant A asserts that a man jumped off the bridge and participant B corrects that statement.

(4.1) A: *A man jumped off the bridge.*
B: *No, he was pushed.* Strawson (1952)

The correction can be paraphrased as "The man (whom you are talking about) did not jump off the bridge, he was pushed (and that's how he got off)." Thus, the first snapshot, made by participant A, shows a man who jumps off a given bridge, at some past time; the second snapshot, made by participant B, shows the same man and bridge at the same past time, but now he is pushed off. The discourse meaning can be compared to the film composed of the snapshots, in other words, the meaning of the discourse is in the *change* from one snapshot to another.

[5]Following Hamblin (1971) *commitments* may be thought of 'as if-beliefs' or, in our own terms, 'public beliefs'. A participant's commitments need not be consistent with his private beliefs, but the other participants will assume this is so as long as there is no evidence to the contrary. We view commitments as a particular kind of attitude.

[6]Within the dynamic semantic paradigm discourse contributions are viewed as update operations. The meaning of a discourse contribution is defined as the change it brings about in the context of interpretation. Formulating it somewhat differently, a participant in a conversation who knows the meaning of a discourse contribution is aware of its update effect, that is, he knows the consequences of accepting its (possibly underspecified) semantic content in the context of interpretation. See Muskens et al. (1997) for an overview.

What we aim at accounting for in our discourse model is the update effect of discourse contributions at the level of semantic contents. For example, the update effect of the correction in (4.1) is to change a description of reality in which a man jumped off the bridge into a description of reality in which that man was pushed off the bridge (and did not jump). The grammar will ensure that if B's utterance is analysed as a correction of A's statement, its semantic content is interpreted as a *rival* description of the situation already described by A. The grammar will predict that accepting the semantic content of the correction in this discourse context involves rejecting the semantic content of the preceding utterance, or part of it.

It remains as a challenge for further research to extend the grammar so as to account as well for the update effect of discourse contributions at the level of the participants' attitudes or commitments. Clearly, every contribution introduces or changes a commitment on the speaker's part, and relates in a manner characteristic for the type of contribution to the commitments already established through the preceding discourse. In the case of the correction in (4.1) B commits himself to a description of reality conflicting with the one participant A is committed to. Presuming A's commitments to be consistent, A can only accept B's correction and at the same time preserve consistency by dropping his commitment to the corrected description of reality.

The commitments of the participants in a conversation relate systematically to the semantic content of their utterances and the types of contributions they make. With respect to assertions, one might say that each participant is committed to what he asserted, or explicitly or implicitly agreed to, but did not retract so far in the discourse (Hamblin 1971). It seems reasonable to suppose that if the language user administrates who makes what kind of contribution at which point in the discourse he is able to infer who is committed to what at any moment in the conversation. Discourse descriptions are sufficiently informative to allow such inferences in principle. Thus, an extension of the discourse grammar may be built which regulates these inferences, and which characterises different types of discourse contributions in terms of the requirements they put on the resulting attitudinal context of interpretation.[7]

[7]An extension of this kind would in effect be a 'grammar of discourse contributions'. Remember that discourse contributions are not described in discourse descriptions in LDG, rather they are modeled by *updates* of discourse descriptions. What characterises an update $\Delta + \delta$ of a discourse description Δ is not just the tree structure described by δ, but also the *attachment* of that stree structure relative to the discourse tree structure described by Δ. While the semantic content of a given discourse contribution could be assumed to be represented by the semantic value of the top node of the tree described by δ, what the speaker commits himself to through his contribution crucially depends on the attachment of the update it induces. The following example illustrates this.

(4.2) A: *If the weather is good we go on a cycle tour.*
B: *And we go swimming.*

Alternatively, a modification of the discourse grammar might be considered in which the semantic values of discourse units are composed of commitments, rather than bare sentence contents. The 'snapshots' resulting from discourse contributions might in that case be descriptions of what we would like to call a *joint* commitment space, containing the information to which all the participants are committed and their individual commitments when these are not shared. Discourse meanings would then represent a neutral outsider's perspective, instead of changing dynamically with the last speaker's perspective. The approach, which basically consists in importing discourse contributions to object level, has potentially far-reaching consequences for the interaction with anaphora resolution, and for discourse structure and the categorisation of discourse relations.

In our view, both approaches are worth pursuing. We cannot tell, at this point, whether one will be more effective than the other. One way or another, an extension of the discourse model that integrates participants' attitudes relative to the things they say shall be necessary if we want to cover the interpretation of discourse contributions in all possible dimensions. We do not pursue such coverage here, however. We concentrate on what must be, in any formal treatment of discourse interpretation, at the basis of the semantic theory, namely the manipulation of bare semantic content through linguistic means.

Consider B's contribution. The semantic value of the top node of the tree described by δ would be $\sigma_k \oplus K$, where σ_k is an underspecified value, and K stands for 'we go swimming'. The most obvious attachment of δ to Δ would be one in which σ_k would stand for 'we go on a cycle tour', and the tree described by δ would form the main clause of the conditional structure. Given that attachment, participant B does not commit himself to 'we go on a cycle tour and we go swimming' by making his contribution, but to 'if the weather is good, we go on a cycle tour and we go swimming'. In other words, the embedding structure matters and the commitment must be inferred nonlocally.

Likewise, the attachment of the update induced by a given discourse contribution determines the content of the appropriateness conditions the contribution imposes on the commitments or attitudes of the participants established through the preceding discourse. (For example, a confirmation may be assumed to trigger a condition that what it confirms must be something which the hearer is already committed to or is proposing to commit himself to.) Thus, the challenge in specifying a 'grammar of discourse contributions' is to account systematically for the communication of semantic content from (object level) discourse descriptions to (metalevel) discourse contributions and vice versa, while doing justice to the individual character of each type of discourse contribution. A very simple approach towards inferring commitments from discourse descriptions would be to suppose that after each contribution the speaker is committed to the discourse meaning resulting from his contribution, while the hearer is still committed to the discourse meaning resulting from his own most recent contribution to the discourse. This, however, would not support an independent characterisation of different types of discourse contributions in terms of their preconditions and update effect on the attitudinal tier of the context of interpretation.

4.3 Nonmonotonic Discourse Contributions

In the preceding section, discourse meaning was defined as $B \oplus \sigma_r$, the update of the implicit background B with the semantic content of the discourse σ_r. Discourse meanings function as contexts of interpretation for discourse contributions. The *meaning* of a *discourse contribution* is its update effect on the context of interpretation established through the preceding discourse. We use $(B \oplus \sigma_r)^{\Delta^{rcl}}$ to indicate the context of interpretation determined by a right-closed discourse description Δ^{rcl}. For a discourse contribution resulting in an update $\Delta + \delta$, its meaning may abstractly be represented as a function from the context of interpretation before processing the contribution to the context of interpretation after processing it: $(B \oplus \sigma_r)^{\Delta^{rcl}} \rightarrow (B \oplus \sigma_r)^{(\Delta+\delta)^{rcl}}$.

Discourse contributions may subsequently be classified in terms of properties of this update function. A distinction we want to highlight here is the one between monotonic and nonmontonic discourse contributions. The difference is illustrated by the following two exchanges.

(4.3) A: *A man jumped off the bridge.*
They couldn't stop him.

In (4.3) the second utterance gives some background information about the event described in the first. The speaker says that a man jumped off the bridge, and then adds that they could not stop him. The second contribution does not change the description of reality put forward through the first one, it only expands it. This is different in the following discourse.

(4.4) A: *A man jumped off the bridge.*
B: *No, he was pushed.*

The second contribution in (4.4), discussed as (4.1) in the preceding section, is a correction. In contrast to the one in (4.3) it *revises* the picture of reality provided through the first contribution: it no longer follows that the man jumped off the bridge. We shall call contributions like the continuation in (4.3) *monotonic*, and contributions such as the correction in (4.4) *nonmonotonic*. The first expand and preserve the context of interpretation, the second induce a revision of the context of interpretation. Using the semantic specification language that will be introduced in the next section, we define these classes of contributions as follows. $B \oplus \sigma_r$, K are of type DRSs and $\models$ denotes an entailment relation on DRSs.[8]

[8]Our semantic specification language CDRT$^+$ is introduced in the next section and fully specified in appendix A.

Monotonic contribution A discourse contribution resulting in an update $\Delta+\delta$ of a discourse description Δ is *monotonic* iff

$\forall K$ if $\Delta^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$ then $(\Delta + \delta)^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$.

Nonmonotonic contribution A discourse contribution resulting in an update $\Delta + \delta$ of a discourse description Δ is *nonmonotonic* iff

$\exists K$ such that $\Delta^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$ and $(\Delta + \delta)^{rcl} \not\models (B \oplus \sigma_r \mid\!\sim K)$.

This says that a discourse contribution resulting in an update $\Delta+\delta$ is monotonic if $(B \oplus \sigma_r)^{\Delta+\delta^{rcl}}$ entails all statements entailed by $(B \oplus \sigma_r)^{\Delta^{rcl}}$, and nonmonotonic if not. In the latter case, there are statements which follow from $(B \oplus \sigma_r)^{\Delta^{rcl}}$ but which do not follow from $(B \oplus \sigma_r)^{\Delta+\delta^{rcl}}$.

Crucially, when we call a discourse contribution nonmonotonic, we are referring to its update effect on the context of interpretation or discourse meaning established by the relevant discourse context. This must carefully be distinguished from its the update effect on the discourse context as such, as represented by a discourse description. Like all other discourse contributions, nonmonotonic ones result in *monotonic* incrementations of the discourse description: one or more logical statements are added to the description and no statement is removed from it.[9] The updates induced by corrections or denials attach to the discourse description like all other updates and do not result in backtracking or a redo of the discourse description. This is indeed as it should be: corrections do not change the interpretation of a previous statement,[10] they only provide conflicting information on the semantic tier of the discourse context. They do not seem to come with a heavier processing load than monotonic discourse contributions.

Interestingly, the monotonicity and nonmonotonicity property are defined in terms of a notion of discourse meaning or context of interpretation which is an underspecified object in our discourse model. At the level of the verifying tree models of a discourse description, this corresponds to a collection of potential semantic values, namely one for each of the verifying tree models of a discourse description. Thus, the meaning of a discourse contribution can be viewed as a relation between a set of potential input contexts of interpretation and a set of potential output contexts of interpretation. An in-depth investigation of the properties of this relation is outside the scope of this dissertation. Due to the interaction with underspecification and ambiguity it is a complex matter. We can, however say something about the distinction between monotonic and nonmonotonic discourse contributions.

[9] The one exception to this is the right-closure assumption, which is hypothetically added to the discourse description after each incrementation step, and dropped when the discourse continues.

[10] Although they may induce the local accommodation of presuppositions which, on the basis of the interpreting agents' preferences, were previously globally accommodated.

To simplify matters, let us consider only discourse contributions which neither introduce new ambiguities nor resolve any of the existing ones in the input discourse description. Furthermore, let us suppose that in these cases there is a one-to-one correspondence between the set of potential input contexts of interpretation and the set of potential output contexts of interpretation, such that intuitively speaking, for every input context there is a *successor* output context which is most similar to it and only differs from it in the update effect of the contribution.[11]

Now it can be seen that monotonic discourse contribution induces a monotonic update on *all* of its potential input contexts of interpretation. So each successor context entails all statements entailed by its input context. Similarly, a nonmonotonic discourse contribution induces a nonmonotonic update on all of its potential input contexts of interpretation, in other words, *no* successor context entails all statements entailed by its input context.

This characterisation of nonmonotonicity covers any contribution which does not preserve the context of interpretation in one way or another. It ranges not just over obvious cases such as denials and corrections, but also over various types of qualifications which arise from stating exceptions or alternatives to a statement, or adding conditions to it. In principle a third category of contributions can be thought of, which might be called 'partially nonmonotonic'. Contributions of this kind would induce updates such that some input contexts are entailed by their successor but others are not. It is an interesting question whether, in terms of conventional classifications of discourse contributions, there is a class which is characterised by this property. Objections or counterarguments, which in a conversation can either be granted and result in a nonmonotonic update, or be explained away and result in a monotonic one, might be examples of this. On the other hand, it is not clear how this property should be defined at the level of the underspecified discourse meanings. So on purely formal grounds it is predicted that this is not a distinctive property in the characterisation of discourse contributions.

4.4 Context Sensitive Compositional Semantics

Discourse contributions relate in many ways to the discourse context, the actual situation in which the conversation takes place, and the world knowledge of the participants. By defining the meaning of a discourse contribution as its update effect on the discourse meaning resulting from the preceding discourse we have placed context dependence of interpretation at the heart of our semantic theory.

[11] We intend to convey only an intuitive understanding of the properties involved; the successor relation, and the similarity relation on which it builds, are not spelled out in any formal detail here. For a formulation of 'similarity' as a primitive relation on (sets of) possible worlds, see for example Lewis (1973).

We can now take a closer look at the linguistic parameters that constrain the relation of a discourse contribution to its context of interpretation.

We will distinguish between *global* requirements and *local* conditions. The first are general conditions on discourse meanings or contexts of interpretation, which must be satisfied after each discourse contribution, irrespective of the type of contribution. In LDG they are formulated as general descriptions or semantic axioms. The well-known requirements that discourse contributions preserve consistency and induce informative updates of the context of interpretation can be formulated as global constraints in the grammar (see section 4.4.5).

Local conditions are constraints on so-called local contexts, induced by specific classes of linguistic elements, such as anaphora, presupposition triggers, discourse connectives and discourse relations. In line with Karttunen (1974), the local context of a linguistic element shall be the contextual semantic information that is accessible to it given the logical structure of the discourse embedding the sentence in which it occurs. All nodes in a discourse tree structure shall be decorated with local contexts (section 4.4.2) and context sensitive constraints shall be integrated in the lexicon of the grammar (section 4.4.3 and 4.4.4). By incorporating constraints on local contexts in the lexical descriptions of context sensitive elements in the language, a unified and declarative treatment of anaphora resolution, presupposition satisfaction, and appropriateness conditions of discourse operators is obtained. Local contexts and context sensitive constraints figure as an additional parameter in the compositional semantics of discourse; thus the result may be called a context sensitive compositional semantics.

Through the combined force of the global requirements on discourse interpretation and the local constraints triggered by the context sensitive elements occurring in a contribution the resulting discourse meaning is constrained. In general that meaning will still be underspecified. On the basis of his world knowledge and default reasoning capacity a language user may generate preferences over different possible readings and obtain a single preferred reading. Thus, pragmatic reasoning interacts with the language user's syntactic and semantic reasoning in the determination of the discourse meaning. We go through an example of this in section 4.4.6 and section 4.4.7. Readers who prefer to see first how the semantic theory will be employed in a concrete example may want to go on to these sections directly and turn back to the preceding ones afterwards.

In the preceding chapters it was simply assumed that linguistic tree representations come with semantic values; no particulars were given about the semantic representation language employed in the description grammar. In this chapter we must fill in the gap. In fact, we need not look for a semantic representation language which differs very much from what is already present in LDG. Recall that the *description language* of the grammar is classical type logic $\mathcal{L}$, and that $\mathcal{L}$ is used to talk about all aspects of linguistic structures including their semantic values. Hence, it figures as the semantic representation language as well. Part of

$\mathcal{L}$ is specifically tuned to the description of semantic values. In Muskens (2001), this part or sublanguage of $\mathcal{L}$ mimics predicate logic. The technique used is virtually identical to the technique used for obtaining the embedding of DRT in $\mathcal{L}$ known as compositional discourse representation theory or CDRT (Muskens 1996; Muskens 2001). Here, we shall employ a treatment of DRT in type logic which is similar to CDRT but offers a somewhat more fine-grained notion of DRSs (Van Leusen and Muskens 2003).

As before, a Logical Description Grammar (LDG) will be a logical theory $\mathcal{G}$, and whenever we have a discourse description Δ_T stating some simple properties of a discourse T, the combined theory $\mathcal{G}+\Delta_T$ describes the syntax and semantics of T. The models of $\mathcal{G}+\Delta_T$ contain trees decorated with semantic values; a function (or, in fact, family of functions) σ associates values with tree nodes. The description logic talks about natural language semantics *and* natural language syntax. In its models we get semantically annotated trees (this in addition to the usual objects and relations) and in its proof theory we can reason about these. As was argued in the preceding chapters *underspecification* naturally falls out of this view. Since a description $\mathcal{G}+\Delta_T$ can have more than one model, the syntax and semantics of T may remain underspecified. In particular, it can be shown that, with the right set-up, certain stages in the reasoning process that takes $\mathcal{G}+\Delta_T$ as its point of departure are closely analogous to the Underspecified Discourse Representation Structures of Reyle (1993). There is no need to postulate a separate structural level of the latter, we get UDRSs for free if the hearer's simultaneous reasoning about the syntax and semantics of an input expression is modeled.

Thus, a powerful formal theory of discourse semantics is integrated in our framework of discourse interpretation. As was noted in the introduction, it is not a technical necessity for LDG to employ DRT as semantic representation language; other choices could be made. This particular choice supports the comparability of our framework with other theories employing (U)DRT and the synthesis of linguistic analyses across these frameworks. In what follows we will assume the reader is familiar with standard DRT as laid down in Kamp and Reyle (1993). As a preliminary to the presentation of our 'context sensitive compositional discourse semantics', in the next section we introduce some basic concepts and notation of finegrained compositional DRT; a full specification of the semantic representation language can be found in Van Leusen and Muskens (2003) and in appendix A.

4.4.1 Finegrained CDRT: Basic Concepts and Notation

Like CDRT, finegrained CDRT is an emulation of DRT in type logic. Semantic representations look like DRS's, but are abbreviations of formulae in type logic mimicking the behaviour of DRT. For example, a DRS condition such as 'w_r : *came-in* u_1' abbreviates $\langle \lambda i.\textit{came-in}(V(u_1)(i), V(w_r)(i)),\ \lambda v[v = u_1 \vee v = w_r]\rangle$ in finegrained CDRT. The variable i ranges over *states* or assignments (type

s), v ranges over registers or discourse referents (type π), and V is a function which delivers the individual that occupies a given register in a given state. A *condition* is a pair of a set of states (the states in which the condition holds) and a set of registers (those that are free in the condition). In the example, $\lambda i.came\text{-}in(V(u_1)(i), V(w_r)(i))$ denotes the set of states such that in that state the individual occupying the register denoted by u_1 is an individual of which the property *came-in* is true in the world occupying the register denoted by w_r, and $\lambda v[v = u_1 \vee v = w_r]$ is the set of discourse referents that are free in the condition.

An essential characteristic of CDRT is that type logic is used to talk at object level about notions which in standard DRT belong to the metalanguage, such as assignments (verifying embeddings) and variable binding. In CDRT registers and assignments are objects in the model, that is, semantic objects. Thus, some of the structure present in the *representations* of standard DRT is imported to the level of the semantic objects they denote.

In finegrained CDRT this applies to the body of conditions of DRS's as well: DRS's are defined so that the semantic object denoted by a DRS reflects the structure of its body, which in the representation is a collection of independent conditions. 'Boxes',[12] such as $[u_1 \mid w_r : \mathit{came\text{-}in}\, u_1,\ w_r : \mathit{sit\text{-}down}\, u_1\,]$ and $[\,\mid w_r : \mathit{sneeze}\, u_1\,]$, abbreviate formulae of type $\pi t \times (st \times \pi t)t$, that is, the semantic value of a DRS is a pair consisting of a set of discourse referents—the *universe* of the DRS—and a set of *conditions*. We shall generally use τ as a shorthand for $\pi t \times (st \times \pi t)t$. This fine-structure is introduced with a view to the treatment of *corrections* later on. It can be argued that corrections and other nonmonotonic discourse contributions induce a split, or a selective downdate of a DRS, with some conditions disappearing or ending up in the scope of a negation while others are preserved. The language user may sometimes reason about arbitrary subsets of a given DRS and, so, must be able to address conditions in isolation.

Given the above typing of DRSs the following two basic operations on DRS's can easily be defined. The *merge* of a DRS K and a DRS K', denoted by $K \oplus K'$, is the DRS resulting from taking the union of the universes, and of the sets of conditions of the two DRS's, respectively. Unlike the merge operation ';' in Muskens (2001), '$\oplus$' is a symmetric operation, in the style of Zeevat (1989). The *inclusion* of a DRS K and a DRS K', denoted by $K \sqsubseteq K'$, holds if the universe of K is included in the universe of K' and the set of conditions of K is part of the set of conditions of K'. For example, we will have that

$$[u_1 \mid w_r\colon \mathit{sneeze}\, u_2\,] \sqsubseteq [u_1\ u_2 \mid w_r : \mathit{came\text{-}in}\, u_1,\ w_r\colon \mathit{sneeze}\, u_2\,]$$

but not

$$[u_3 \mid w_r\colon \mathit{sneeze}\, u_3\,] \sqsubseteq [u_2 \mid w_r\colon \mathit{sneeze}\, u_2\,]$$

[12] We employ representations with brackets instead of the boxes of standard DRT: DRS's are of the form $[\ldots \mid \ldots]$. To the left of the | sign is the *universe* or set of declared discourse referents of the DRS, to the right its body of *conditions*.

We shall sometimes want to require that a DRS is proper, or contains no free ('unbound') discourse referents. The *free* referents of a DRS K are the ones that are free in some condition of K but do not belong to the universe of K. We shall use *proper*(K) to denote that K has no free referents.

Furthermore, we shall distinguish three kinds of registers, generated by u, o, and w, functions from nodes to registers. The semantic axioms of the grammar ensure that fresh registers come with each node in a tree structure. The values of u will be associated with *new* referents, while the values of o are referents that belong to the background B of the discourse. The difference between o- and u-generated referents will help to distinguish names from pronouns and other anaphoric elements.

The values of w are registers that can store *worlds*. While states correspond to assignments and registers correspond to variables in predicate logic, the notion of a possible world of course corresponds to the technical notion of a *model*. Possible worlds play an essential role in defining a notion of entailment at object level, a matter to which we turn below. In this dissertation we will only use $w(r)$, the w value of the root.[13] The occupant of $w(r)$ is the 'actual world' or world of evaluation. It will be ensured that w_r is always present in the universes of the 'local contexts' of nodes, and indeed in the global context, or *background* B. A notational convention that we find useful is to write arguments of u, o, and w as subscripts, for instance u_3 instead of $u(3)$, w_r instead of $w(r)$ etc.

In our treatment of context sensitive aspects of interpretation we will want to put constraints on contexts and typically these constraints can take two forms. One possibility is to require that a context *contains* certain material. For example, for lexical nodes k that carry the proper name 'Pedro', we will demand that $[o_k \mid w_r{:}\, \mathit{Pedro}\ o_k] \sqsubseteq B$. Another possibility is that a DRS K is required to *follow from* a context. We will see an illustration of this in chapter 6, where presuppositions shall be required to be satisfied in their local contexts. In this chapter, it shall be exemplified by the implementation of appropriateness constraints associated with discourse relations and discourse operators. For this purpose, a notion of entailment *at the level of the describing logic* is made available. It is obtained by quantification over worlds: an entailment holds if the conclusion is true in all worlds in which the premises are true. We employ $K \mid\!\!\approx K'$ to indicate that a DRS K entails a DRS K'.

All of the notions introduced above are defined in appendix A. This also contains a truth-preserving translation from our CDRT-sublanguage to predicate logic, and a definition of truth stated in terms of this. An alternative dynamic treatment along the lines of Groenendijk and Stokhof (1991) is briefly discussed. In the course of the chapter some further notation will be introduced, to allow us to talk with some precision about aspects of the semantics whose full formal treat-

[13] But in extensions containing modal operators it would be natural to use other values of w as well.

ment awaits future research. We shall use $\stackrel{\pi}{\sim}$ to indicate object level entailment given the application of common sense reasoning or default inference to the left argument. Furthermore, we shall introduce a function ϱ, from nodes to registers, the occupants of which are local reference points. They shall figure in statements constraining the spatio-temporal and mereological structure of discourse. This presupposes an event semantic extension of the semantic representation language.

4.4.2 Decorating Tree Structures with Local Contexts

We set out to provide a discourse semantics that can deal with context sensitivity in general. In particular, we promised to provide a unified and declarative semantics of anaphoric or presuppositional elements in the language. The treatment that we have in mind consists in integrating *local contexts* and constraints on them in LDG. More specifically, the elementary tree structures in the lexicon of the grammar shall be outfitted with local contexts, so that every node in a linguistic tree structure comes with a local context. Local contexts may then serve as a parameter in the compositional semantics of context sensitive elements. Different classes of anaphoric or presuppositional elements can be characterised in terms of the specific constraints on local contexts they trigger.

The notion of a local context derives from Karttunen (1974). Karttunen proposes to account for the projection of presuppositions triggered within complex sentences in terms of the 'satisfaction' of the presupposition relative to the relevant context of interpretation, as established by the complex sentence. Contexts are modeled as sets of logical forms, and satisfaction comes down to entailment. In subsequent work, this proposal has been taken up and formulated in terms of updates on sets of possible worlds and definedness conditions on contexts, see for instance Heim (1983). Here, we are interested in the basic idea, namely that the context of interpretation for a presupposition triggered by a given, simple sentence in a discourse, is a function of the logical composition of the discourse preceding and embedding the sentence.

As Karttunen says, "in compound sentences, the initial context is incremented in a left to right fashion giving for each constituent sentence a local context that must satisfy its presuppositions". Local contexts of embedded sentences are defined recursively in terms of the logical form of the compound sentence that is the embedder. Karttunen provides such definitions for compound sentences of the form 'If A then B', 'A and B', 'A or B', and sentences with a sentential subject or object. Though Karttunen talks about presuppositions only, the approach carries over to constraints triggered by context sensitive elements in general.

The manner in which local contexts are constructed in our discourse grammar is essentially the same; the main difference is that in LDG, local contexts are DRS's rather than sets of logical forms, and of course the available discourse syntactic structure is employed in defining them. For each node k in a tree structure,

its local context $\Gamma(k)$ is computed top-down from the local context of its mother node (if it has one) and the semantic values of its (left) sister nodes. This is regulated by means of the elementary structures in the lexicon and the semantic axiom $\mathcal{A}18$ below. The axiom sets the local context of the root node $\Gamma(r)$ to the implicit background or global context of the discourse B (a type τ constant), and it says that nodes of a syntactic category other than S inherit the local context of their mother node. Here $k_1 \lhd k_2$ abbreviates $k_1 \lhd^+ k_2 \wedge \neg\exists k[k_1 \lhd^+ k \wedge k \lhd^+ k_2]$, that is, $\lhd$ denotes the immediate dominance relation.

$$\mathcal{A}18\ \ \Gamma(r) = B \ \wedge\ \forall k_1 k_2\,[[k_1 \lhd k_2 \wedge \ell(k_2) \neq s] \rightarrow \Gamma(k_2) = \Gamma(k_1)]$$

The computation of the local contexts of S-nodes other than the root is captured in the lexicon of the grammar. To illustrate this, consider the following lexical descriptions. They describe the syntax and semantics of the sentence connectives *and* and *if*. As in chapter 2, we employ pictures to represent lexical descriptions. The conventions that are followed in our graphic representations are the same as before. In addition, the semantic value $\sigma(k)$ of a node k may be written under it and we typically write σ_k for $\sigma(k)$, and Γ_k for $\Gamma(k)$.

(4.5)

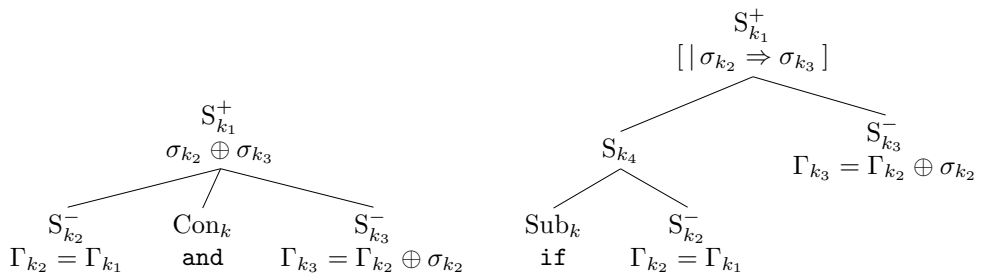

In each case, node k_1 will receive its local context from some negatively marked node it can be identified with, when the lexical description of the construction features in a sentence or discourse description. Thus, Γ_{k_1} shall be the local context resulting from the preceding discourse. In the lexical description of *and*, Γ_{k_2}, the local context of the left conjunct, is identified with Γ_{k_1}, and the local context of the right conjunct, Γ_{k_3}, is an update of this with σ_{k_2}, the semantic value of the left conjunct. Similarly, in the lexical description of *if*, the local context of the antecedent of the conditional is identified with Γ_{k_1}, and the local context of the consequent is an update of this with the semantic value of the antecedent. All of this reflects Karttunen's treatment of compound sentences composed with *and* or *if* directly. His theory of presupposition satisfaction may be implemented in LDG by requiring presuppositions to be *entailed* by the local contexts of the linguistic elements that trigger them.

Following Karttunen (1974), the lexical description of the connective *or* may be given as in (4.6), or as in (4.7).

(4.6)

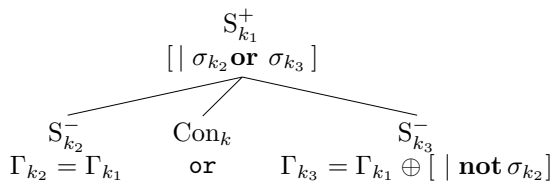

In (4.6), the local context of the second disjunct is the local context resulting from the preceding discourse updated with the negation of the semantic value of the left disjunct.

(4.7)

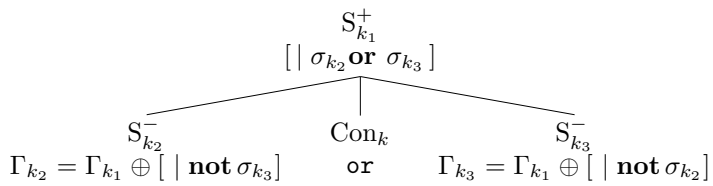

In (4.7), the local context of the each disjunct is the local context resulting from the preceding discourse updated with the negation of the semantic value of the other disjunct. Karttunen discusses both possibilities. Each comes with its own—desirable or undesirable—consequences for the theory of presupposition. What matters to us, however, is not so much the particular choices to be made for each connective in the language, but rather, the general approach, the integration of local contexts as such in the semantics of LDG.

All elementary tree structures anchored by (explicit) discourse operators and (implicit) discourse relations in the lexicon shall be decorated with local contexts in this manner, so that, with the help of axiom $\mathcal{A}18$, the local context of each node in a discourse tree structure can be composed top-down. It may be observed that given the composition of local contexts as laid down in the grammar, the local context of every node in a discourse parse tree contains the implicit background B. Furthermore, the local contexts of discourse units or S-nodes, like their semantic values, are all DRSs.

Local contexts relate to the notion of *accessibility* in standard DRT in the sense that whenever a discourse referent is accessible from an a pronoun in DRT, in our discourse grammar, that referent is in the universe of the local context of the node at which the pronoun sits, and, hence, can bind or resolve it. What is contained in the local context of a pronoun includes at least the information one comes across in standard DRT when walking the accessibility path from the pronoun to the main DRS disregarding information that may have been added through subsequent incrementation steps.[14] For any given node in a discourse tree

[14]In some cases the Karttunen-style local context contains more information than what would be the information collected on the accessibility path in standard DRT. For instance, when a discourse contains a disjunction, the Karttunen style local context of the right argument of the disjunction contains the negation of the semantic value of the left disjunct. In standart DRT, however, neither the semantic content of the left disjunct, nor its negation is on the accessibility

we shall say that its local context represents the information which is *semantically accessible* at that node.

4.4.3 Lexical Descriptions of some Basic Word Classes

Let us take a closer look at the lexicon, and zoom in on some basic word classes. We will discuss the semantics of the determiner *a*, and of (instances of) verbs, common nouns, names, and pronouns. The point of interest is, of course, the treatment of context sensitivity in the lexicon. As discussed in the preceding section, elementary tree structures are decorated with local contexts. Here we shall see that anaphoric elements such as names and pronouns come, in addition, with constraints on their local contexts.

As in the preceding chapters, there will be *classifying* descriptions and *elementary tree* descriptions in the lexicon. Of these, the elementary tree descriptions carry most information, but the classifying descriptions play a useful role in connecting elementary tree descriptions with open class words. We speak of a *lexical description* when, for a given lexical item or discourse relation, the classifying and elementary tree descriptions are combined.[15]

In (4.8) below classifying descriptions for the open class words *Pedro, has,* and *mule* are given. (4.8a) states that whenever a node k carries the word *has*, it must be classified as a transitive verb (tv) and its semantics $\sigma^{\pi(\pi\tau)}(k)$ is $\lambda v'\lambda v$ $[\ |\ w_r{:}\ v\ \mathit{has}\ v']$. (If α is a type, we will let σ^{α} denote a function from tree nodes to objects of type α.) The description for *mule* in (4.8b) is similar, but the word is classified as a common noun (cn) and its semantic value is of type $\pi\tau$. Lastly, the classifying description for *Pedro* in (4.8c) sets the semantics to a type $\pi\tau$ property as well but classifies the element as a proper name (pn). Differences between common nouns and proper names will become apparent when we consider the elementary tree descriptions they select.

(4.8)a. $\forall k[\textsc{has}(k) \rightarrow (tv(k) \wedge \sigma^{\pi(\pi\tau)}(k) = \lambda v'\lambda v[\ |\ w_r{:}\ v\ \mathit{has}\ v'])]$

b. $\forall k[\textsc{mule}(k) \rightarrow (cn(k) \wedge \sigma^{\pi\tau}(k) = \lambda v[\ |\ w_r{:}\ \mathit{mule}\ v])]$

c. $\forall k[\textsc{pedro}(k) \rightarrow (pn(k) \wedge \sigma^{\pi\tau}_k = \lambda v[\ |\ w_r{:}\ \mathit{Pedro}\ v])]$

As before, classifying descriptions will be associated with elementary tree descriptions via the wordclasses or syntactic categories they select (tv, cn, pn, etc.). Elementary descriptions specify the syntactic structure and semantic features lexemes of a given word class provide, or subcategorize for. Like classifying descriptions, they are formulated in type logic. As in chapter 3, however, we prefer to use graphical notation for elementary descriptions. The official representations

path from the right disjunct upto the main DRS. If wanted, the local contexts in LDG can also be made to correspond exactly to what is on the accessiblity path in standard DRT.

[15] In Van Leusen and Muskens (2003) both classifying and elementary descriptions are called lexical descriptions.

of elementary tree descriptions can always be reconstructed from the more user-friendly pictorial representations.

The picture in (4.9) is a graphic representation of the elementary tree description for elements of word class *tv*, that is, for transitive verbs like *has*, in (4.8a). It says that whenever a tree contains a transitive verb, it must have certain other properties as well. The transitive verb must be labeled 'V', it must have an immediately dominating node labeled 'VP', which in turn is dominated by a (possibly identical) node labeled 'VP', whose left sister is a node labeled 'DP', etc. The description also contains anchoring information and semantic information. For brevity, superscripts on σ indicating the type of a semantic value are often dropped.

(4.9) *tv:*

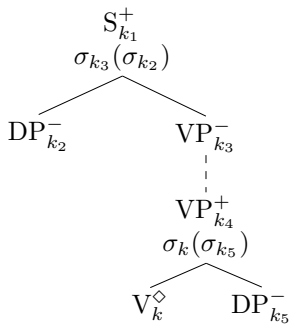

The description in (4.9) covers the whole class of transitive verbs. If it is conjoined with the information that k carries the lexeme *has* and with the classifying description (4.8a) we get the lexical description of the word, graphically represented by (4.10). Note that the semantics of VP_{k_4} now has a more complete specification.

(4.10)

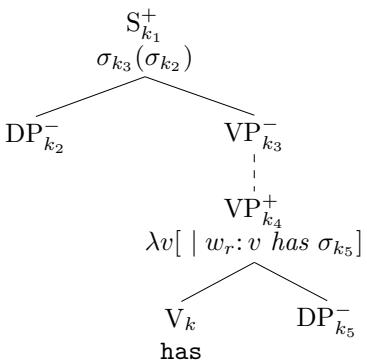

For common nouns (*cn*) and proper names (*pn*) we will have the following elementary tree descriptions.

(4.11) *cn:*

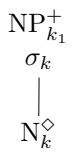

pn:

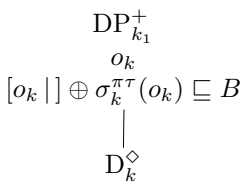

The elementary description for common nouns sets the semantic value of the maximal projection (the NP) to the semantic value of the noun. The elementary description for proper names is of a slightly different nature. Firstly, it sets the semantics of the DP to the discourse referent o_k; this referent will figure as an argument of a predicate (verbal or otherwise) when the DP node is identified with an open slot in some other elementary tree. Secondly, it introduces a requirement which enforces that the implicit background of the discourse B contains the DRS $[o_k \mid] \oplus \sigma_k^{\pi\tau}(o_k)$, where $\sigma^{\pi\tau}$ is the property specified in the classifying description of the proper name. This in fact implements the usual DRT requirement that discourse referents connected with names must be present in the universe of the main DRS and that the descriptive material connected with such discourse referents must likewise be available globally.[16]

Combining the elementary descriptions in (4.11) with the classifying descriptions in (4.8c,b) we obtain the lexical descriptions of *mule* and *Pedro*.

(4.12)

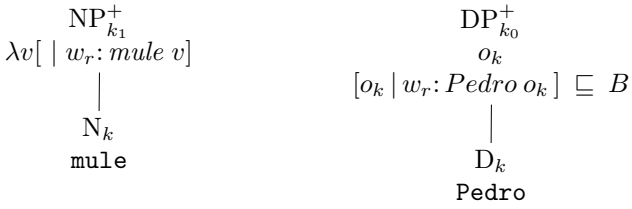

For *closed class* words it is often expedient to represent the individual lexical items directly by means of a lexical description. Here, for example, is an entry for the indefinite *a*.

(4.13)

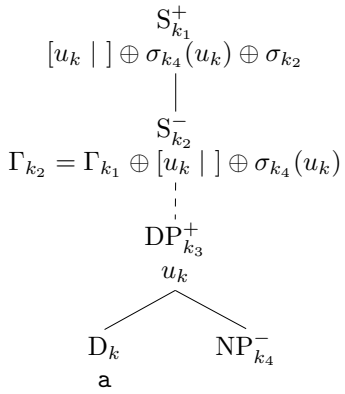

The idea here is that the DP that is projected from the indefinite translates as a discourse referent u_k and that the declaration of that discourse referent

[16] The *o*-generated discourse markers belong to the implicit background or global context of the discourse. A *referential* interpretation of proper names may be obtained by adopting a truth definition for DRSs which ensures that discourse markers present in the universe of the implicit background receive their interpretation directly from the state of evaluation i. This will be discussed in section 4.4.7.

$[u_k \mid]$ and its restriction $\sigma_{k_4}(u_k)$ are quantified-in at some higher S level. The semantic value of the DP, u_k, functions exactly like a trace and in the process of semantic composition it is bound by the quantifier higher up in the structure. S_{k_1} should be compared to the place where adjunction of the DP takes place after Quantifier Raising in generative grammar.[17] Note that Γ_{k_2}, the local context of the lower S node, is set to be the merge of the local context of the higher S node and the 'restrictor' material of the semantics of that higher node in (4.13). Unresolved pronouns will seek a referent in this local context, and presuppositions may likewise resolve to it.

The determiners *no* and *every* may similarly receive lexical descriptions as in (4.14).

(4.14)

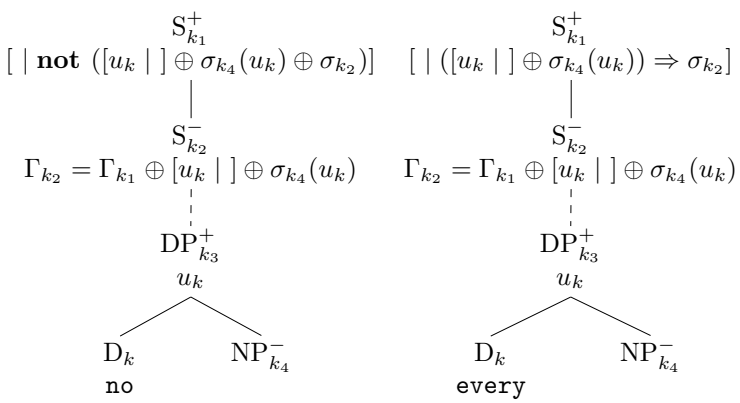

Finally, here is a lexical description of the pronoun *he*.

(4.15)

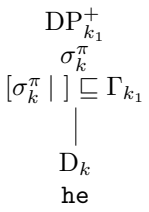

Pronouns, like names, come with a semantic value and an additional constraint on their local context. The semantic value of the DP is set to σ^π_k, a register generated at the terminal anchor, node k. The statement $[\sigma^\pi_k \mid] \sqsubseteq \Gamma_{k_1}$ requires this value to be present in the universe of the local context of the DP. Unlike discourse referents introduced by indefinites the semantic values of pronouns are underspecified. The statement $[\sigma^\pi_k \mid] \sqsubseteq \Gamma_{k_1}$ enforces their identification with a semantically accessible discourse referent, thus implementing the binding condition on pronouns familiar from DRT.

The lexicon of the grammar is now sufficiently specified to generate descriptions of sentences like 'Pedro has a mule', 'He has Pedro', etc. The two context

[17] This treatment of quantifiers in general is motivated in Muskens (2001). See also Cimiano and Reyle (2005) for some critical discussion.

sensitive, or anaphoric elements in this tiny fragment of English are characterised semantically not just in terms of their semantic values but also of particular constraints on local contexts. In the next section we will zoom in on the semantics of the elements that 'glue' sentences together, that is, discourse connectives and discourse relations. As may be expected, these shall be treated as context sensitive elements in the lexicon as well. After that, in section 4.4.5, we shall talk about the implementation of two nonlexical, or *global* semantic constraints on discourse interpretation, namely consistency and informativity. The interaction of context sensitive constraints and semantic values in the composition of discourse meanings will be discussed when we go through the interpretation of an example discourse in sections 4.4.6 and 4.4.7.

4.4.4 Discourse Relations, Connectives, and Adverbials

In interpreting discourse, a language user constructs a coherent whole on the basis of what is explicitly stated. Every utterance is taken to relate in some more or less specific way to what was previously conveyed in the discourse. Whether or not that relation is explicitly described through discourse connectives and other discourse operators, the language user will infer that it is there. Consider

(4.16) A: *John was late. He had a flat tire.*

The most preferred reading of this discourse is one in which there is a causal relation linking the second statement to the first: John was late because he had a flat tire. We assume or infer that the second statement is an explanation of the first. The speaker might have used the connective 'because' to convey this explicitly. Very often, however, as in (4.16), coherence relations are partly or completely underspecified. The task of a discourse grammar is to explain that, despite this, discourse is interpreted as a coherent whole.

As we saw in chapter 3, following common practice in discourse theory, we propose to account for coherence in discourse in terms of a collection of *discourse relations*, representing or modeling the relations of coherence that glue together the utterances in a discourse. In our treatment, discourse relations are implicit elements in the lexicon of the grammar. They characterise coherence relations only partially in the sense that they do not describe the impact of coherence relations on the participants' commitments, intentions or attitudes. Apart from this however, in principle every discourse relation fully describes the content-parameters of a given coherence relation, so that every coherence relation whose semantic-pragmatic impact is purely at the level of content is fully characterised by a single discourse relation. Optionally, in line with Asher and Lascarides (2003), it may be assumed that coherence relations stand for speech act types, or *types* of contributions.[18]

[18]Ones intuitive understanding of what speech acts and coherence relations are may point in the direction of a one-to-one correspondence between coherence relations and speech act types,

Note that we distinguish between discourse relations, which are implicit elements in the lexicon of the grammar, and coherence relations, the objects in reality they model or describe. Coherence relations can also be explicitly described through the lexemes in an utterance. In fact, they are always at least partially determined through the presence of discourse connectives and other discourse operators, and through the tense, aspect, mood and modality of the sentence uttered, and its information structure. The linguistic properties of the sentence constrain, but do not necessarily uniquely determine a coherence relation. Even when explicit connectives are present coherence relations can be underspecified. As in Knott and Mellish (1996) we account for all this by letting discourse operators describe *properties* (or in their terminology: *features*) of coherence relations rather than full coherence relations. This immediately also explains that discourse operators can in principle be 'stacked' or combined (Webber et al. 1999).

Discourse relations Discourse relations are not phonologically expressed but, like lexemes, they anchor elementary tree structures, and each relation comes with its own description in the lexicon of the grammar. The axioms of the grammar ensure that discourse relations are in complementary distribution with (two-place) discourse connectives: when the one is present, the other is excluded. Furthermore, if a tree representation of a sentence is linked to the discourse context through a discourse relation, discourse syntax allows the presence of maximally one such element.

Discourse relations are semantically characterised through the combination of their semantic value and a particular set of context sensitive constraints or *appropriateness conditions*. The latter constrain the relation between the semantic value of the newly processed discourse unit which they link to the discourse context and its local context. In addition, through the selection of a coordinating or a subordinating elementary structure, discourse relations have a particular impact on the topical structure of the discourse. (Though the discourse topical structure is a parameter of discourse interpretation still to be filled in in our grammar).

To illustrate, here is an entry for the discourse relation *continuation*, of word-class *rel*. *Rel* is a closed class, the elements of which are represented by type νt predicates or labels such as *continuation*, *explanation*, *elaboration*, and *correction*. For easy reference, we decorate conditions on local contexts with labels such as [**inf**], [**cnq**], and [**inc**], abbreviating 'informativity', 'consequence', and 'incompatibility' respectively; the labels are not part of the descriptions, however.

---

but it is a moot point what the implications of that assumption will be for speech act theory. Since in our grammar discourse relations model coherence relations, and discourse units are composed of discourse relations and clauses or complex discourse units, it would follow that speech acts can be composed of (embedded) speech acts, and we would end up with a typically linguistic perspective on speech act theory. An alternative approach would be to assume that more than one coherence relation may contribute to a single speech act type.

(4.17)

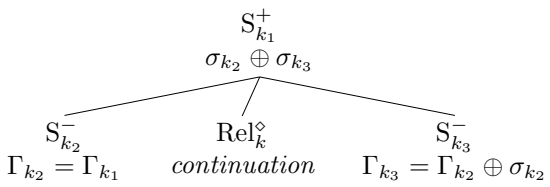

The *continuation* relation[19] is a coordinating relation which combines a contextually given discourse unit (at node k_2) with a subsequent clause or discourse unit (at node k_3). The local context of the contextual argument (Γ_{k_2}) of the continuation is set to the local context of the whole node (Γ_{k_1}). The local context of the right argument (Γ_{k_3}) is set to that of the left argument updated with the left argument's semantic value. Semantically, a continuation results in the merge of the semantic values of its arguments. Optionally, an appropriateness condition [**inf**] $\Gamma_{k_3} \not\models \sigma_{k_3}$ could be included in (4.17). This would require the *informativity* (non-entailment) of the right argument relative to its local context. Informativity conditions are discussed in more detail in section 4.4.5.

The semantic value and appropriateness conditions of a discourse relation may involve any of the parameters that constrain discourse meaning: logical relations, spatio-temporal relations, part-whole relations, aspectual distinctions, modality and causality. As we haven't made available an eventuality semantic extension of finegrained CDRT, not all of this can currently be expressed in our semantic representation language. To be able to talk about these parameters nonetheless, we will introduce some schematic notation.

For example, consider the following entry in the lexicon, for the discourse relation *explanation*. A state of affairs A may be assumed to explain a state of affairs B if in the context of interpretation A causes B. Hence, *explanation* comes with the appropriateness condition [**cau**]. The condition requires that the semantic value of the right argument of the discourse relation (the explanation) 'causes' the semantic value of the left argument (the thing to be explained) given the local context of the right argument. The predicate **cau** in the semantics denotes a primitive relation of causation. This is purely notation; the predicate may be seen as a place-holder for a theory of causation which is yet to be specified.

[19] Adopted from Asher and Lascarides (2003). The continuation relation is a semantically not very constraining discourse relation which resembles what is called a 'list' structure in the Linguistic Discourse Model. As far as we can see a continuation cannot be used in an out-of-the-blue context at the beginning of a discourse, in other words, it usually plays a subsidiary role with respect to some other, more specific or constraining discourse relation. This could easily be implemented in LDG by adding constraints to the lexical description of the continuation relation.

(4.18)

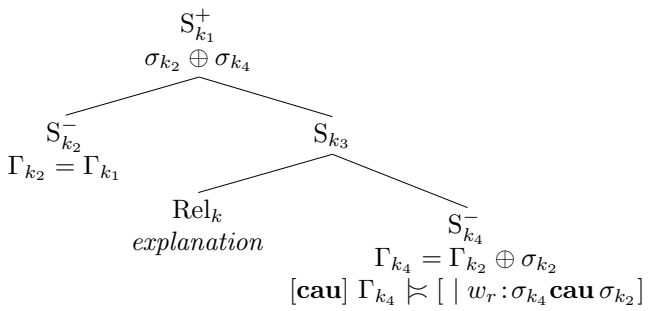

In our discussion of example (4.16) we claimed that the coherence relation linking 'He had a flat tire' to 'John was late' is a relation of explanation. We expect our framework of discourse interpretation to predict that when interpreting the discourse in (4.16) the hearer infers the discourse units representing the two sentences are linked through the implicit discourse relation *explanation*. It may be observed however that given the descriptions in (4.17) and (4.18) whenever *explanation* can be inferred, the relation *continuation* can be inferred as well. The first is more specific than the second. Though *continuation* does not enforce a causal relation, it is compatible with it. Thus, it is possible to infer *continuation* while the discourse gets exactly the same interpretation as with *explanation*, resulting in a void prediction of ambiguity or underspecification.

For reasons of semantic hygiene we want to avoid this. We will assume that whenever the language user's preference system generates a preference for a certain reading, the most specific discourse relation allowing that reading is selected. This situation must be distinguished from the one in which there are two possible, distinct readings of a discourse, one of them more specific than the other. For instance, 'He had a flat tire' could be interpreted as a neutral continuation of 'John was late', where it is undecided whether there is a causal relation or not. Then there is a real ambiguity between a *continuation* reading and an *explanation* reading of the discourse. According to our intuitions about this particular example, the second reading is preferred. This might be explained on the assumption that the language user's preference system contains a general preference for more specific or semantically stronger readings (and hence, discourse relations). The selection of *maximally coherent* readings in general may be explained from the interaction with other preferences, such as for instance a preference for local attachment, or a preference for a most normal or uncontroversial specification of the implicit background B.

The lexicon of the grammar may also be taken to include a description of discourse relation *result*, which is, as it were, the inverse of *explanation*.

(4.19)

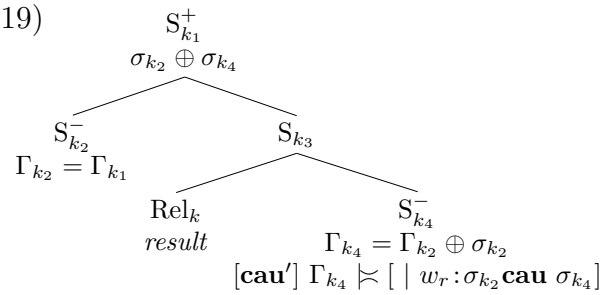

This relation applies to examples like the following, where the most preferred interpretation is one in which we infer that the eventuality described by the first statement causes the eventuality described by the second, given the context of interpretation.

(4.20) *John pushed Alice. She fell.*

For some discourse relations, their main characterising property is their impact on the spatio-temporal and mereological relations within discourse. This applies, for example, to elaborations and narrations. *Narrations* are coordinating relations which induce a relation of temporal sequence between their arguments. If both arguments describe events, the event described by the right argument must temporally succeed the one described by the left argument. If the arguments describe states or events, the state or event described by the right argument should not temporally precede the one described by the left argument. A lexical description generalising over these cases this might look as follows.[20]

(4.21)

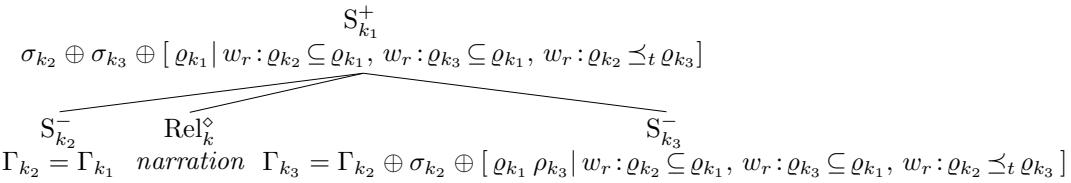

The CDRT$^+$ language employed in (4.21) is furnished with some notation signifying eventualities and temporal relations. This is basically cosmetics, and builds on the assumption that an eventuality semantic extension of the semantic representation language can be made available. We employ a ϱ-operator (of type $\nu\pi$) which generates what we propose to call *local reference points* for each S-node in an elementary description. The values of ϱ are registers that can store eventualities. We assume that local reference points govern the temporal referential

[20] This treatment of *narration* allows that its arguments convey either states or events. One may, of course, distinguish cases in which both arguments express an event from event-state, or state-event sequences, and introduce different discourse relations for each case. This results in a more fine-grained classification of discourse relations.

structure of discourse, and can be related in the appropriate way to each other and to temporal discourse referents or discourse referents of type eventuality generated from the lexical input. A condition of the form $w_r : \varrho_k \preceq_t \varrho_{k'}$ says that a local reference point ϱ_k temporally precedes or starts at the same time as a local reference point $\varrho_{k'}$ in world w_r; $w_r : \varrho_k \subseteq \varrho'_k$ says that ϱ_k is part of (or equal to) ϱ'_k in world w_r.

Semantically, *narration* produces the merge of the semantic values of its arguments and introduces a local reference point ϱ_{k_1} which is required to contain the local reference points of the arguments ϱ_{k_2} and ϱ_{k_3}. These local reference points are temporally ordered through $\varrho_{k_2} \preceq_t \varrho_{k_3}$. The effect of this should be that within some implicitly understood time span given by ϱ_1 two (complex) eventualities took place, each in their own time span constrained by ϱ_{k_2} and ϱ_{k_3}, where the first does not start later than the second. On the assumption that constraints on the spatio-temporal referential structure are part of the basic compositional semantics of discourse relations they are included directly in the semantic value of the resulting discourse unit k_1 and in the local context of the right argument of the discourse relation, Γ_{k_3}. Indeed, we assume that *all* discourse relations to a greater or lesser extent constrain the spatio-temporal and mereological structure of discourse. The corresponding constraints will be represented in our descriptions of discourse relations only when it is relevant to the discussion, however.

The discourse relation *elaboration* is usually taken to be a subordinating relation, whose main semantic characteristic is that the semantic value of the subordinated right argument zooms in on the semantic value of the head of the construction, represented by the left argument. This can be implemented in our grammar by requiring that the local reference point of the right argument of the elaboration is part of the local reference point of the left argument. This is what the following lexical description says.

(4.22)

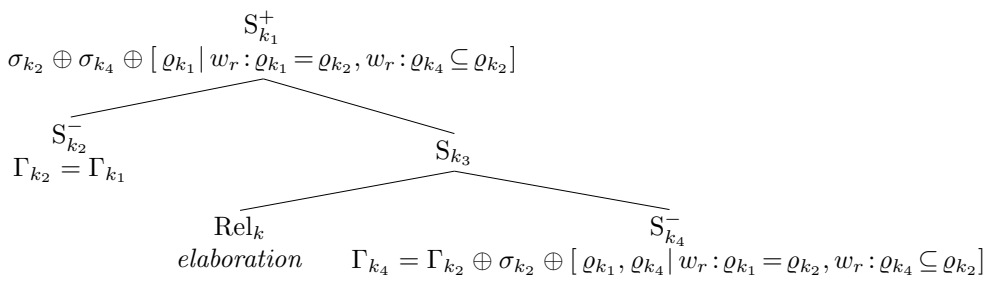

Instead of the rather space-consuming pictures of descriptions such as (4.21) and (4.22) we will generally employ more succinct representations like the ones in (4.23) below. Here the conditions on the spatio-temporal and mereological relations are listed under the semantic value of the top node, and labeled [**lrp**], which stands for **l**ocal **r**eference **p**oint. The pictures in (4.23) simply abbreviate the ones in (4.21) and (4.22) and represent exactly the same lexical descriptions.

(4.23)

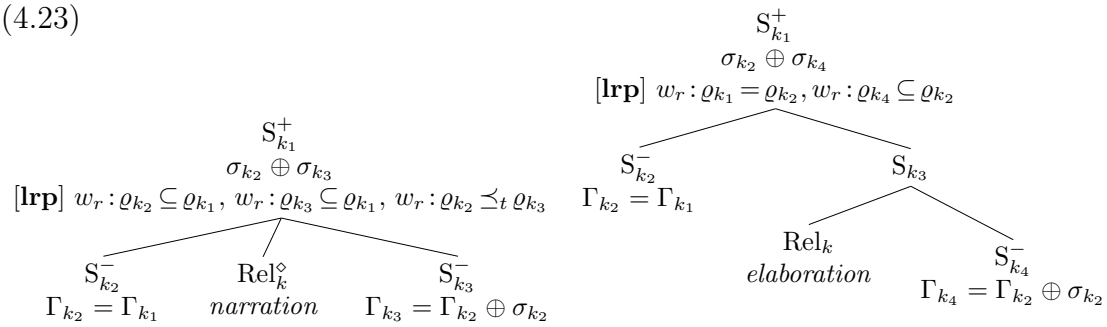

A final distinction we would like to introduce here is the one between veridical and nonveridical discourse relations. Veridical discourse relations or discourse operators imply the truth of the semantic values of their arguments. Building on Zwarts (1995) and Asher and Lascarides (2003) we shall call a discourse relation or a discourse connective *veridical* if the semantic value it results in (that is, the semantic value of the topnode of the elementary structure it anchors) entails the semantic values of its arguments in arbitrary contexts of interpretation, and *nonveridical* if it does not. The discourse relations discussed in this section so far are veridical. This is obviously the case when only the 'merge' operator is employed in composing the output semantic value of a discourse relation or connective. Nonveridical connectives are 'if', 'or' and 'not'; an example of a nonveridical discourse relation is *denial*, the semantic impact of which might be defined as in the following lexical description.

(4.24)

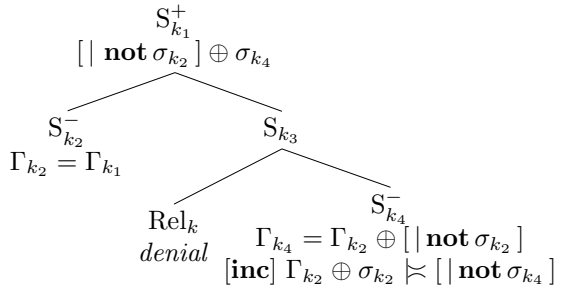

The semantic value of the top node is the negation of the semantic value of the contextual argument merged with the semantic value of the right argument. Thus, the contextual information that is rejected through the denial is captured in the scope of a negation.[21] Given this output value, *denial* is a nonveridical discourse relation.

Veridicality could be viewed as a local monotonicity property of discourse relations. In section 4.3, monotonicity was defined as a *global* property of discourse

[21] As observed in various places (van der Sandt 1991; Asher 1995; van der Sandt and Maier 2003; van Leusen 2004) a treatment like this both over- and undergenerates. Our treatment is slightly improved in chapter 6, where it is made sure that, through the interaction with local accommodation, presuppositional material can be captured in the scope of denials. Further refinements are necessary, however.

contributions. A discourse contribution resulting in an update $\Delta{+}\delta$ is monotonic if and only if its semantic update effect preserves the input discourse meaning: $(B \oplus \sigma_r)^{\Delta+\delta^{rcl}} \models (B \oplus \sigma_r)^{\Delta^{rcl}}$, and it is nonmonotonic if not. Corrections and denials are nonmonotonic discourse contributions. It is not necessarily the case, however, that nonmonotonic discourse contributions employ nonveridical discourse relations. Incrementing a discourse description with a nonveridical discourse relation or connective embedded in the argument of another nonveridical one may result in a monotonic update, while the use of a veridical discourse relation or connective embedded in the argument of a nonveridical one may induce a nonmonotonic update.

As can be seen from (4.24), discourse relation *denial* comes with an appropriateness condition [**inc**]: $\Gamma_{k_2} \oplus \sigma_{k_2} \models [\,|\, \mathbf{not}\, \sigma_{k_4}\,]$. As is argued in Asher (1995) and van Leusen (2004) it may be assumed that what makes a denial appropriate in its context of interpretation is that the semantic content of the denial is inconsistent with it. The *incompatibility* condition [**inc**] expresses this. Note that the appropriateness condition is defined relative to the 'nonrevised' local context $\Gamma_{k_2} \oplus \sigma_{k_2}$, while Γ_{k_4}, the local context which will be percolated downwards, is revised: the 'culprit' information σ_{k_2} is replaced by its negation. Thus, the semantic impact of the denial is taken into account in the local contexts of discourse constituents embedded in the right argument of the denial.

Sometimes a *veridical* discourse relation or discourse operator comes with a kind of incompatibility requirement too. This typically applies to discourse operators like 'but', 'however', and 'nonetheless'. All of these deny or reject some default expectation supported by the context of interpretation. The following example illustrates this.

(4.25) *John insulted me. He is a nice guy, however.*

If John insulted the speaker, presumably, he is not a nice guy. The second utterance conveys that, contra expectations, John *is* a nice guy. Crucially, the information that John insulted the speaker is not rejected as a result of the second utterance. Discourse contributions like these cannot be handled just like denials and corrections because the information they reject is not the semantic value of their contextual argument. Just adding [**inc**] as an appropriateness constraint would result in an inconsistent discourse meaning, which is incorrect. The only way to handle such cases properly seems to be to allow the semantic representation language to talk about commonsense inference or the 'normally implies' relation.

Following Lakoff (1971) we shall call the relation of coherence involved 'denial of expectation'.[22] We may introduce the discourse relation *denial-exp* in the

[22]Most current treatments of the discourse connectives mentioned analyse the relation involved as one of *contrast* rather than denial of expectation. We view denial of expectation as an instance of contrast. Contrast can be treated as a requirement of incompatibility on the

lexicon of the grammar to model it. The relation is characterised by the appropriateness constraint [**n-xp**]. This features the predicate $\rightsquigarrow$, which stands for a relation of commonsense inference or default implication. The condition [**n-xp**] says that in the local context of the right argument of the discourse relation the semantic value of the left argument normally implies the negation of the semantic value of the right argument.

(4.26)

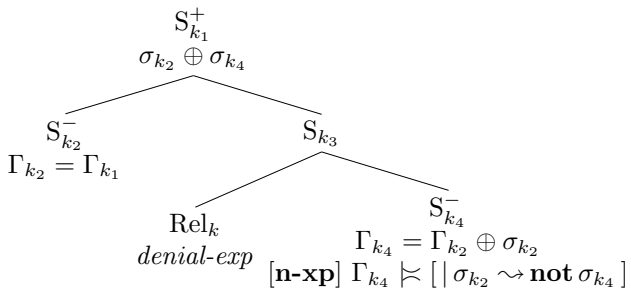

In the case of example (4.25), it may be assumed that the hearer selects the discourse relation *denial-exp* as a suitable link between the two statements. The commonsense rule defined by [**n-xp**] would be instantiated as 'if John insulted me, it normally follows he is not a nice guy'. The appropriateness condition enforces the selection of a local context in which this commonsense rule holds. Thus, in the absence of counter-information, Γ_{k_4} supports the default expectation that John is not a nice guy. We will say that Γ_{k_4} *pragmatically entails* that John is not a nice guy, and employ the notation $\Gamma_{k_4} \overset{\pi}{\models} [\,|\, \mathbf{not}\, \sigma_{k_4}\,]$ to represent this.[23] Crucially, the discourse meaning resulting from the second utterance does *not* support the default expectation, even though [**n-xp**] is satisfied. Since the speaker asserted that John is a nice guy, the commonsense rule does not 'fire': $B \oplus \sigma_r \not\overset{\pi}{\models} [\,|\, \mathbf{not}\, \sigma_{k_4}\,]$. The negative expectation is cancelled, as it were.

Discourse operators Like discourse relations, discourse connectives, adverbials and other discourse operators are context sensitive elements. They come with appropriateness conditions or presuppositions characterising their relation to the context of interpretation. These conditions describe properties of coherence relations; they constrain the coherence relation that links the internal argument of the discourse operator its the discourse context. The appropriateness conditions attached to a discourse operator need not uniquely determine a coherence relation, but may underspecify it. Furthermore, discourse operators can be com-

local context, defined in terms of potential semantic alternatives of the internal argument of the connective. As we have not provided a definition of 'semantic alternatives' we keep to the simpler treatment in terms of denial of expectation here.

[23] For a common sense rule or default implication $K' \rightsquigarrow K''$ we would have that if $\Gamma \models K' \rightsquigarrow K''$ and $\Gamma \models K'$ then $\Gamma \overset{\pi}{\models} K''$. Similarly for a causal implication K' **cau** K''. Furthermore we would have that if $\Gamma \models K$ then $\Gamma \overset{\pi}{\models} K$, but not vice versa.

bined in a single utterance as long as their syntactic and semantic contraints do not clash and the total of their appropriateness conditions can be satisfied.

When implicit discourse relations are inferred, their characterising properties must be consistent with the constraints triggered by the discourse operators that are explicitly present. To illustrate this, consider the following lexical descriptions, characterising the discourse connective 'so' and the adverbial 'consequently'.

(4.27)

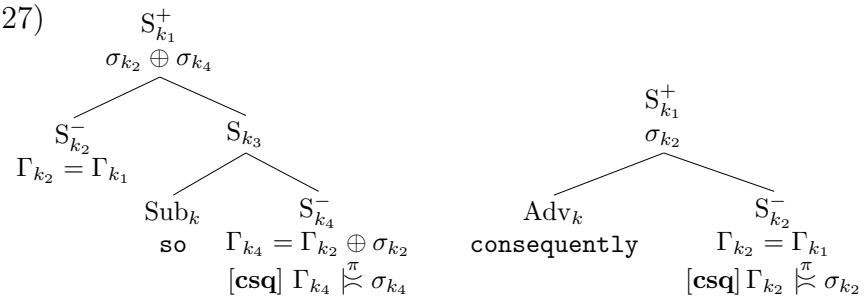

Both elements come with the appropriateness condition [**csq**] ('consequence'). It requires that the semantic value of the internal argument σ_{k_4} follows from its local context Γ_{k_4}, given the potential use of commonsense reasoning. We follow Webber et al. (2003) in treating discourse adverbials as one-place, rather than two-place operators. Webber c.s. argue that discourse adverbials behave anaphorically in that they admit the crossing of structural dependencies, and have a wider range of options with respect to their external argument than do coordinating and subordinating connectives. In our grammar the appropriateness conditions express the anaphoric or presuppositional relation to the context.

As we saw in chapter 2, in the absence of an explicitly conveyed structural link to the preceding discourse, the grammar predicts the presence of an implicit discourse relation, which anchors the structural link. Obviously, the discourse relation that is selected must be compatible with the semantic value and the appropriateness conditions attached to the discourse adverbial. Consider

(4.28) *John didn't leave the party in time. Consequently, he missed the train.*

The adverbial 'consequently' requires that 'John missed the train is a probable consequence in a context in which John did not leave the party in time. This is obviously so if the fact that John did not leave the party in time caused him to miss the train. Hence, the implicit discourse relation *result* is a suitable candidate. Furthermore, on the basis of the preferences mentioned earlier in this section, implicit discourse relations that are compatible with the adverbial but not as specific as *result* will be excluded. The discourse description in (4.29) shows the resulting structure and constraints.

(4.29)

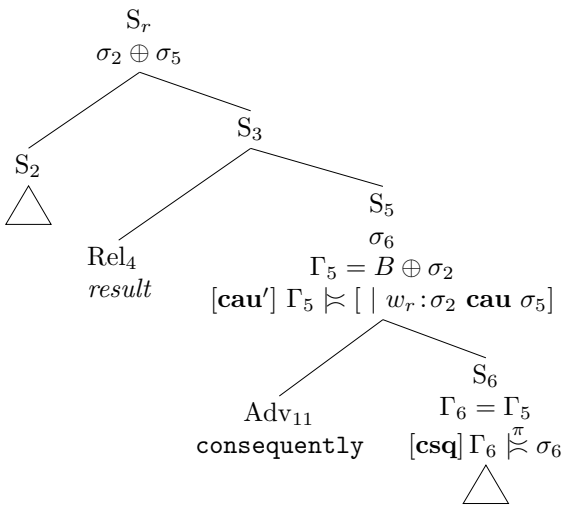

In general, discourse adverbials and other discourse operators constrain, but do not necessarily uniquely determine the implicit discourse relation (and hence, the coherence relation) that links their internal argument to the discourse context. On the basis of his world knowledge and preference system, the language user may strengthen the interpretation of the utterance, and thereby select any of the coherence relations that fits the requirements of the connective. As the following example shows, even when no discourse relation is inferred, because an explicit discourse connective is present, coherence relations can be underspecified.

(4.30) a. *Peter plays basketball and Jack plays badminton.*
b. *John came in and took off his hat.*
c. *John didn't leave the party in time and missed the train.*
d. *John missed the train and didn't leave the party in time.*

In (4.30a), the connective 'and' simply conjugates two statements: it tells us that the eventualities (states) described by its arguments hold. In (4.30b) and (4.30c) something more is inferred. In (4.30b) we assume that John took off his hat *after* coming in, in (4.30c) that John missed the train *because* he didn't leave the party in time. The interpretations obtained correspond to having a narration relation, and a relation of result, respectively. Interestingly, in (4.30d) the inference that John missed the train because he did not leave the party in time seems to be blocked. The most preferred reading of (4.30d) is one in which the conjuncts sum up unrelated facts or eventualities, which corresponds to a relation of continuation.

Given these observations we propose that the compositional semantics of the connective 'and' is simply the merge of the semantic values of its arguments, as in (4.31) below, while the interaction with the language user's world knowledge

and preference system should explain the strengthening of the readings occurring with the connective.

(4.31)

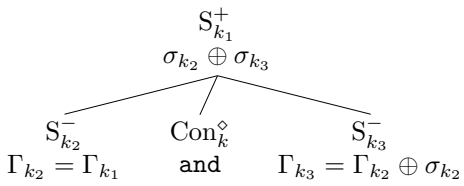

The observations made relative to (4.30) could be explained on the assumption that the language user's preference system includes a preference for assigning a temporal order to a pair of consecutive discourse units that reflects their sequential order in the discourse. This would explain why in (4.30b) we assume that John took off his hat after coming in, instead of before, why in (4.30c) we assume that John's being in a state of not leaving the party in time started before he missed the train.[24] and why, in (4.30d) we do *not* infer that John missed the train because he did not leave the party in time. The not leaving in time would temporally precede the missing of the train, and this would violate the preference. Finally, in (4.30a), the preference would be overruled by a general preference for uncontroversial or nonexceptional background assumptions.

Things are more complicated than that, however. Consider the following minimal pair, which repeats (4.30d) as (4.32a).

(4.32) a. *John missed the train and didn't leave in time.*
b. *John missed the train. He didn't leave in time.*

In (4.32b) it is completely unproblematic to infer that John missed the train because he did not leave in time, in other words, there is a relation of explanation. Why doesn't the preference for inferring a temporal order that follows text order prevent this reading? The explanation can be maintained if we accept that the preference is attached to the connective 'and' specifically. But this means that preferences may be set differently for different elements in the lexicon, which seems unattractive. Further research is needed, to find out how the preference system can be fit to the subtleties of the linguistic data.

As in Webber and Joshi (1998), our discourse model predicts a difference between aspects of interpretation supported by the compositional semantics of discourse, and aspects of interpretation induced by the world knowledge and preferences of the language user. The latter are cancellable, but the first are not. Crucially, context sensitive constraints are part of the compositional semantics of lexical elements and discourse relations in our treatment. When an element in a sentence comes with a context sensitive constraint, the constraint must be

[24]The independent preference for chosing a more specific reading when there is a choice would account for the additional causal link in (4.30c).

satisfied and its implications are predicted to be uncancellable. This holds even when the constraint in question talks about default inferences or commonsense reasoning, as is the case with [**n-xp**], or [**csq**] in (4.19). Indeed, as the following example shows, if the condition is part of the discourse description because it is triggered by a lexical element in the discourse, it cannot be cancelled.

(4.33) *John didn't leave in time. Consequently, he missed the train.*
\# *As the train was late he missed it for another reason however.*

In contrast, if a default conclusion results from pragmatic strengthening of a reading on the basis of world knowledge and preferences over potential readings, it can be cancelled. Indeed, although the cancellations in (4.34) may be unexpected, they are not unfelicitous.

(4.34) a. *John didn't leave in time and he missed the train.*
As the train was late he missed it for another reason however.
b. *John didn't leave in time. He missed the train.*
As the train was late he missed it for another reason however.

Summary and discussion Summarising, in this section it was proposed that implicit discourse relations, discourse connectives and other discourse operators are context sensitive elements. They come with appropriateness conditions which describe properties of coherence relations. Discourse relations are implicit elements in the lexicon which stand for full coherence relations. Discourse connectives and adverbials partially describe coherence relations through their appropriateness conditions. Discourse operators and discourse relations can be combined as long as the total of their context sensitive conditions can be satisfied. Despite the presence of explicit discourse operators, coherence relations may be underspecified. The interplay with the language user's reservoir of world knowledge and preference system explains the selection of most preferred readings or coherence relations.

Clearly, the specification of the preference system, and the reasoning involved in generating preferences over alternative readings of discourse, is an area for further research. It is a common assumption in cognitive linguistic and discourse semantics that language users strive to maximise discourse coherence in their interpretations. Preferences should be set so that possible readings that are intuitively more coherent than others will be preferred. The parameters involved in this are diverse and often gradient. Asher and Lascarides (2003) mention, among more, the number of anaphoric expressions whose antecedents are resolved, and the specificity of a common topic induced by a rhetorical relation (more in general, the quality of the rhetorical relation). The preference for semantically more specific interpretations introduced in this section could be added to this. Syntactic preferences, such as the preference for low attachment to the discourse

structure proposed in Gardent and Webber (1998) may be taken to contribute to the maximisation of coherence too. As we have seen, unlike what follows from the compositional semantics of a discourse, conclusions based on preferences are cancellable in the presence of new information. Thus, cancellability may be employed as a criterion in deciding whether aspects of discourse interpretation should be captured by means of the preference system or not.

In our choice of discourse relations and their semantic characterisation we more or less followed Asher and Lascarides (2003). The fragment presented here is far from complete, however, both in the sense that many discourse relations have been left undiscussed, and in the sense that certain aspects of their characterisation have not been specified. Notably the interaction between discourse relations and the information structure of their arguments awaits further research. What we have been concerned with is not so much full descriptive coverage. Rather, we have tried to show the potential of LDG as a description language for discourse theory.

As a typically grammatical approach towards coherence, our treatment is closer to TAG for discourse (Webber and Joshi 1998; Gardent and Webber 1998) and the Linguistic Discourse Model (Scha and Polanyi 1988). The LDG formalism does not enforce a particular taxonomy or categorisation of discourse relations. The class of discourse relations in the lexicon is only required to be a closed class. Optionally, an integrity constraint might be adopted which says that the characterising properties of discourse relations do not differ from those that characterise the explicit lexemes of the language, so that as a class, implicit discourse relations are a natural extension within the lexicon of the language. What is accepted as a suitable set of discourse relations must come out of practical experience with what is needed to account for different types of discourse, and be based on extended investigation of the possibilities of use of explicit discourse operators in the language.[25]

Given the crucial role of the preference system in accounting for the selection of most preferred readings and corresponding discourse relations, it can even be argued that once a fully worked out preference system is available, the class of discourse relations could be reduced to a single implicit discourse relation of 'adjunction' or 'update'. This relation is semantically underspecified, but the lan-

[25]See for example (Knott 1996; Knott et al. 2001; Marcu 1997; Marcu 2000; Sporleder and Lascarides 2007). We like to set a pointer in particular to Sporleder and Lascarides, who test the use of automatically labelled examples to classify rhetorical relations in text, even when no explicit discourse operator is present. Their results suggest that training on this type of data, where explicit discoure operators are as it were replaced by rhetorical relations, may not be such a good strategy, as classifier models trained in this manner do not seem to generalise very well to data containing no explicit operators. They suggest that a reason for this may be that examples with and without explicit discourse operators are too dissimilar linguistically. This fits in nicely with our own hunch, that implictness suits some coherence relations better than others. Furthermore, it confims the idea that it is not just discourse connectives, cues, and adverbials, that constrain or determine coherence relations.

guage user infers a suitable semantics on the basis of his preferences.[26] The specific discourse relations that were introduced in the lexicon of the grammar before can then be viewed as conventionalised representations of the usual outcome of the inference process. They do not in themselves explain *why* a particular interpretation is inferred, but they do tell us something about *what* interpretations are normally inferred. As such they play an essential role in the investigation of the preference system, even though eventually they can be eliminated from the grammar.

4.4.5 Global Consistency and Informativity Conditions

It is commonly assumed in discourse semantics that for a contribution to a conversation to be 'felicitous' or 'appropriately used' relative to the context established by that conversation, various requirements must be satisfied. In the wake of Austin (1962) and Searle (1969), a theory of discourse interpretation may focus on the requirements or felicity conditions that are *specific* to particular types of contributions or speech acts, and which characterise them as, for example, promises, assertions, questions, warnings, or objections. Building on Grice (1975), Grice (1978), it may also aim at specifying *general* requirements characterising cooperative conversation.

In the LDG framework of discourse interpretation both types of requirements can be modeled. General requirements characterising cooperative conversation can be implemented as *global* conditions on discourse meanings. And felicity conditions that are specific to particular types of contributions may be implemented as *local* constraints, in other words, constraints on local contexts, triggered by specific linguistic elements. In the latter case, it is tacitly assumed that there is a systematic relation between discourse relations and the types of contributions occurring in conversations, so that the constraints attached to a given discourse relation reflect (part of) the felicity conditions characterising a discourse contribution of the same name.

Here we discuss the treatment in LDG of two general conditions, descending from the Gricean maxims of Quality and Quantity respectively. We shall refer to them as the consistency and the informativity condition. These two are widely accepted as a minimal set of felicity conditions for discourse aimed at information exchange. They figure as acceptability conditions on updates in presupposition theory (Stalnaker 1978; van der Sandt 1988) and in update semantics (Veltman 1996). The consistency condition says that discourse contributions must preserve

[26] This seems to converge with the approach taken towards discourse relations in Webber et al. (2003). Their LTAG for discourse contains an elementary structure anchored by punctuation, the semantics of which is inferred on the basis of the interaction of compositional discourse semantics, anaphor resolution, and inference triggered by adjacency. A precursor of the latter component figures in Webber et al. (1999) as 'inference based on world knowledge, usage conventions, etc., which can make defeasible contributions to discourse interpretation'. This would correspond to our preference system.

the consistency of the context of interpretation. The informativity condition requires that every contribution provides new information: the update it induces must 'make a difference' and change the context of interpretation.

Both requirements can in principle be formulated as global constraints on discourse meanings in LDG. However, while the consistency requirement can straightforwardly be stated at object level in the description language, the informativity requirement refers to different stages of the discourse description and can only be stated at metalevel.

Global Consistency $B \oplus \sigma_r \not\models \bot$

Global Informativity For a discourse contribution updating a discourse description Δ to $\Delta + \delta$, there must be a DRS K such that

$$\Delta^{rcl} \not\models (B \oplus \sigma_r \mid\!\approx K) \text{ and } (\Delta + \delta)^{rcl} \models (B \oplus \sigma_r \mid\!\approx K).$$

Global Consistency prevents that a discourse contribution results in an inconsistent discourse meaning. Global Informativity ensures that every discourse contribution results in a discourse meaning which is at least partly new: there must be information K entailed by the output discourse meaning which is not entailed by the input discourse meaning. We will adopt Global Consistency as a semantic axiom of the grammar. With respect to the other requirement, however, it is less clear that it should be accepted. Different types of contributions come with their own felicity or appropriateness conditions, and in some cases these are not compatible with Global Informativity. With certain caveats and restrictions, though, the requirement might still be adopted.

Consistency Global Consistency interacts in an interesting way with the monotonicity property of discourse contributions. It enforces that contributions which add information that is inconsistent with what was previously stated or presupposed in the discourse result in nonmonotonic updates. If a nonmonotonic reading is excluded the contribution is predicted to be infelicitous. Consider the examples in (4.35) and (4.36).

(4.35) A: *It is raining.*
B: *No, it isn't.*

If B's utterance in (4.35) would be interpreted as a monotonic contribution, the resulting discourse meaning would be inconsistent and Global Consistency would be violated. Hence, B's utterance can only be interpreted felicitously as a nonmonotonic discourse contribution, replacing 'it is raining' by 'it is not raining'. Then, the resulting discourse meaning is consistent, so that Global Consistency is satisfied. In the case of (4.35) such a reading is available: B's utterance can be interpreted as a denial of A's statement. Now consider

(4.36) A: *It is raining and it is not raining.*

In contrast, in (4.36) a nonmonotonic reading is not available, due to the presence of the connective 'and'. The semantic value that 'and' produces is the merge of the semantic values of its arguments and 'merge' is a *veridical* operator, that is, it preserves the semantic values of its arguments. Since the statements which 'and' connects contradict each other, its use results in a violation of Global Consistency. The language user's grammar will fail to come up with a reading for this discourse, and (4.36) is predicted to be semantically anomalous, or ungrammatical on semantic grounds.[27]

Notice that the Global Consistency condition is formulated independently of the language user's reservoir of world knowledge $\mathcal{K}$. When interpreting a discourse, he will evaluate discourse meanings relative to $\mathcal{K}$. Obviously, checking the consistency of the discourse meaning with respect to $\mathcal{K}$ will be part of the evaluation. In general, he will have a preference for readings of the discourse that are consistent with his world knowledge. As the exchange in (4.35) shows, however, it may very well occur that a statement made by another participant in the conversation is inconsistent with his world knowledge, without being semantically anomalous. The language user may evaluate a statement like that relative to his world knowledge $\mathcal{K}$ by computing the *revision* of $\mathcal{K}$ that would be induced if he were to accept the discourse meaning resulting from the statement while preserving consistency. His preference system would tell him which of the alternative possible revisions to prefer. Discourse meanings are not in general required to be consistent with the hearer's world knowledge.

Finally, observe that the effect of Global Consistency could also be obtained by integrating local conditions on consistency in the lexical descriptions of discourse operators and discourse relations. To illustrate, the following would be a description of the connective *and* including such a local consistency requirement:

(4.37)

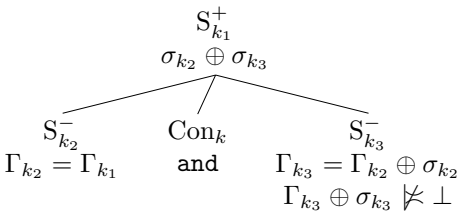

$\Gamma_{k_3} \oplus \sigma_{k_3} \nvdash \bot$ is a condition on the local context of the right argument of the connective updated with the semantic value of that argument. Post-conditions of

[27] It may be observed that in some situations, for instance when explaining to a person that he cannot be right, or when reasoning ad absurdum, the aim of the argument is expressly to establish inconsistency. Exchanges or arguments of that kind should not be predicted to violate Global Consistency and hence to be semantically anomalous; they should somehow escape the effect of the condition. This might be done by making sure that the inconsistency they create is captured in the scope of a negation in their output semantic value.

this kind could be added to all relevant elements in the lexicon. Since the outcome of *every* discourse contribution must be a consistent discourse meaning, however, we prefer to treat consistency as a global condition.

In some cases, it can be argued that particular types of discourse contributions induce consistency or inconsistency requirements on *input* discourse meanings. This applies, for example, to denials and corrections, which may be claimed to be felicitous only if an inconsistency with the context of interpretation can be established. As was proposed in the preceding section, these can be handled by means of pre-conditions on local contexts triggered by discourse relations and connectives. This is perfectly compatible with Global Consistency.

Informativity Global Informativity is a more worrisome notion, both in terms of the intuitive judgements it is supposed to cover, and in terms of its interaction with other constraints in the grammar formalism. The fact that the requirement can only be formulated at a metalevel in our grammar seems ominous: we are considering a notion which in a sense falls outside the scope of our discourse grammar, and which more properly belongs to what we previously called a separate grammar of discourse contributions.

It is not clear whether Global Informativity applies to each and every type of discourse contribution. Different types of contributions come with their own set of appropriateness constraints, and the interaction of these constraints with Global Informativity is a matter for careful consideration. In some cases the individual appropriateness conditions of a contribution make Global Informativity superfluous as an independent constraint. This goes, for example, for corrections and denials. We assume that what makes them appropriate is that their semantic content is incompatible with information present in the context of interpretation, and that they induce nonmonotonic updates of the context of interpretation. Global Informativity is necessarily satisfied. The informativity condition does not seem to add anything to the analysis of these contributions.

In other cases the interaction between Global Informativity and the specific appropriateness conditions of a particular type of contribution is more subtle. Consider the following example, in which the second utterance expresses the consequence of some event or state of affairs described in the preceding discourse. Assume the discourse takes place in a context in which nothing was previously stated or presupposed about the occurrence of a strike of the train personnel and it is not known to the hearer that the trains are not running.

(4.38) A: *There's a strike of the train personnel.*
So the trains are not running.

The fact that the trains are not running is, in the eyes of the speaker, a consequence of the fact that there is a strike of the train personnel. Given general knowledge about the relations between strikes, train personnel and trains, under

normal circumstances it can be expected that the trains are not running if there is a strike of the train personnel. A contribution of this kind is felicitous only if what it asserts can be (and is) understood as a default consequence of the context of interpretation. As was proposed in the preceding section, this requirement can be implemented as a local constraint or appropriateness condition attached to the discourse connective *so*, as follows.

(4.39)

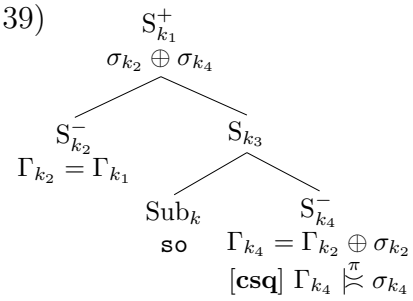

Whereas [**csq**] requires pragmatic entailment of 'the trains are not running' in a local context which corresponds to the discourse meaning described by the input discourse description Δ^{rcl}, Global Informativity requires that adding the same information to that discourse meaning is informative: the discourse meaning described by the output discourse description $(\Delta + \delta)^{rcl}$ should entail something not previously entailed. For the sake of transparency, let us assume that [**csq**] is defined in terms of $\mid\!\sim$, so we have [**csq**] $\Gamma_{k_4} \mid\!\sim \sigma_{k_4}$ instead of [**csq**] $\Gamma_{k_4} \overset{\pi}{\mid\!\sim} \sigma_{k_4}$. Allowing pragmatic entailment in [**csq**] complicates, but does not undermine the argument made here.[28]

We expect a violation of Global Informativity. Surprisingly, however, this does not occur. This is because the local condition [**csq**] constrains only the discourse meaning described by the output description $(\Delta + \delta)^{rcl}$, and not the discourse meaning described by the input description Δ^{rcl}. Remember discourse meanings are underspecified. The update effect of participant A's second utterance is to include the consequence relation in the implicit background B, and eliminate readings of the discourse in which it does not hold. Thus, for σ_{k_4} = 'the trains are not running', we have $\Delta^{rcl} \not\models (B \oplus \sigma_r \mid\!\sim \sigma_{k_4})$ and $(\Delta + \delta)^{rcl} \models (B \oplus \sigma_r \mid\!\sim \sigma_{k_4})$ and Global Informativity is satisfied. The contribution is informative because it teaches the hearer that the trains are not running and that this is a consequence of the strike of the train personnel.

[28] A pragmatic entailment $\Gamma \overset{\pi}{\mid\!\sim} K$ can hold while entailment $\Gamma \mid\!\sim K$ does not. Hence, examples like the one in (4.38) may seem to escape a violation of Global Informativity in case they build on default inferences and commonsense reasoning rules alone. Again, however, [**csq**] applies to the output discourse meaning, rather than the input discourse meaning. If [**csq**] induces accommodation, that explains why GI is not violated. Only if we consider cases in which [**csq**] does not induce accommodation can it be argued that the satisfaction of GI enforces essential use of commonsense reasoning in the satisfaction [**csq**].

In fact, what this means is that, as far as Global Informativity is concerned, a contribution which induces the *accommodation* of information counts as informative.[29] The condition only predicts the infelicity of contributions which do not teach the hearer anything new in any way, that is, cases in which both the presupposed and the asserted material is already established as an entailment of the discourse meaning described by the input discourse description. The problem is that, relative to arbitrary contexts of interpretation, we do not know beforehand what is in B; though contributions headed by *so*, like the one in (4.38), very often induce accommodation, they cannot be assumed to do so generally. Thus, this type of contribution still poses a problem for Global Informativity. This is definitely the case with modus ponens arguments that are spelled out in conversation, as for instance in

(4.40) A: *During a heat-wave, the school is closed.*
We're having a heat-wave right now.
Consequently, the school is closed.

Consider the last contribution in (4.40). The input discourse meaning of this contribution entails that the school is closed. Global Informativity cannot be satisfied and the contribution is incorrectly predicted to be infelicitous. It may tentatively be concluded that the predictions in the case of (4.38) are accidentily correct, or right for the wrong reasons.

We are not the first to observe that modus ponens arguments in discourse flout the informativity requirement. A common way out of this problem is to assume that Global Informativity only applies to discourse contributions or discourse aimed at information exchange. It does not apply to logical reasoning. The question is, of course, how such a restriction could be made concrete in a formal discourse model such as ours. An alternative solution would be to drop Global Informativity altogether and account for the 'standard' cases of informativity by means of *local* informativity constraints only. Local informativity conditions could be attached as appropriateness conditions to some of the discourse relations and discourse operators in the lexicon.

For example, consider the following, extended lexical description for the connective *and*. The condition [**inf**] requires the second conjunct of the construction to be informative relative to the local context established by the first conjunct.

[29]The term *accommodation* was introduced in Lewis (1979) to denote the implicit process of adjustment of the context of interpretation of a statement, when without such adjustment no suitable interpretation can be got. Our underspecification formalism allows us to model accommodation as exactly what it is: the specification of hitherto underspecified implicit background information, triggered by the interpretation of the ongoing discourse.

(4.41)

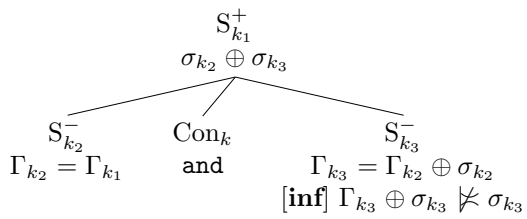

Decorating discourse operators and discourse relations at large with local informativity conditions may easily cause problems, however. As the following example shows, if we actually adopt (4.41), the resulting theory of informativity may well be too strong.

(4.42) A: *There's a strike of the train personnel. Consequently, the trains are not running and there are questions in Parliament.*

Consider participant A's second utterance. We assume the conjunction *and* is in the scope of the discourse operator *consequently*. The adverbial *consequently* triggers the [**csq**] requirement on 'the trains are not running and there are questions in Parliament', while the connective *and* triggers the [**inf**] requirement on 'there are questions in Parliament'. But then, the two appropriateness conditions threaten to contradict each other. Only the fact that [**csq**] is defined in terms of *pragmatic* entailment while [**inf**] is defined in terms of entailment saves the example from a conflict between the constraints. It is questionable, however, whether all contributions headed by *consequently* or *so* necessarily specify default consequences. If they do not, the grammar predicts a conflict where it should not.

Thus, including an informativity condition in the lexical description of the connective *and* may not be a good idea after all. Clearly, if informativity is to be accounted for in terms of local informativity conditions, fine-tuning is necessary. Further research will have to show whether perhaps only some classes of elements in the lexicon, for instance discourse relations, come with local informativity conditions. If that would not bring us closer to getting observationally correct predictions, we might consider implementing local informativity constraints as weak constraints or preferences, rather than hard constraints in the grmmar, as they are now.

Finally, there are types of contributions that are informative in a way that escapes Global Informativity altogether because as it stands our semantics does not talk about participants' commitments, intentions or attitudes. Consider the following exchange, in which the second utterance serves as a confirmation of the first.

(4.43) A: *Freddy is a marvelous guitar player.*
B: *He is.*

Confirmations often just repeat information that is already present in the context of interpretation. At the level of semantic content, they do not necessarily add any new information. Rather, they are informative in the sense that they convey a new speaker's commitment. Global Informativity incorrectly predicts these contributions to be infelicitous, and thus should not be taken to apply to such cases. Which aspects of confirmations can be captured by means of local informativity conditions is, once again, a matter for further investigation.

We conclude that informativity is a much more diversified concept than consistency. Despite the intricacies of the data and the issue at hand, it may tentatively be assumed that Global Informativity captures a characteristic of a certain class of discourse contributions, namely those which act at the level of plain semantic content, and which figure in discourse aimed at information exchange. Given some future extension of the grammar that covers the participants' commitments, intentions and attitudes, Global Informativity might be maintained as a semantic or pragmatic constraint of the discourse theory. Unlike Global Consistency, it applies at metalevel rather than object level in our grammar formalism. Local informativity conditions can be assigned to discourse relations or discourse operators when Global Informativity should not suffice to account for the linguistic observations, or when fine-tuning of the theory of informativity is desired. Care should be taken, however, that the resulting set of constraints is not too strong. In chapter 5 our discussion of informativity and consistency requirements is continued.

4.4.6 Inferring Semantics from Discourse Descriptions

In the preceding sections, we talked about discourse meanings and context dependence of interpretation. context sensitive elements, such as names, pronouns, discourse relations, and discourse connectives, were introduced in the lexicon, and two global constraints on contexts of interpretation, namely consistency and informativity, were introduced as semantic axioms of the discourse theory. We are finally ready to consider the interaction of all constraints in the analysis of a concrete example.

Consider the interpretation of the following discourse. The procedure is exactly the same as in chapter 3. Upon hearing the sentence

(4.44) *Pedro has a mule.*

the language user will construct a discourse description, on the basis of an input description summing up the lexemes in the sentence as immediately consecutive elements, as in (4.45).

(4.45) $\exists k_0 k_1 k_2 k_3 (k_0 \prec_l k_1 \prec_l k_2 \prec_l k_3 \wedge lex(k_0) \wedge lex(k_1) \wedge lex(k_2) \wedge lex(k_3) \wedge$
$\textsc{pedro}(k_0) \wedge \textsc{has}(k_1) \wedge \textsc{a}(k_2) \wedge \textsc{mule}(k_3))$

This being the first sentence of a discourse, the hearer will know that *Pedro* is the start element of the discourse and conclude that the discourse description is left-closed. As the discourse can be taken to be temporarily finished at this point, the hearer may reason about the discourse description under the assumption that it is right-closed as well. Combining the input description with the appropriate elementary tree descriptions in his lexicon, and taking (fresh) witnesses for the existential quantifiers in the elementary tree descriptions, the resulting discourse description will be something like (4.46).[30]

(4.46) Δ_1^{rcl}: S_r^-, $\Gamma_r = B$

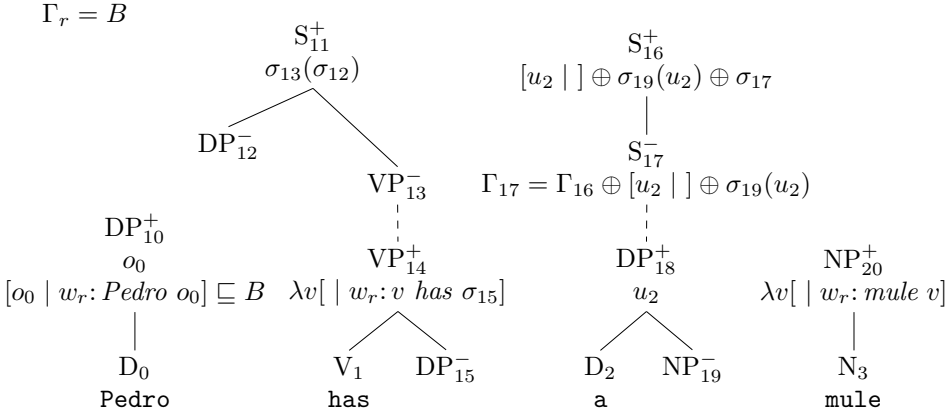

As before, the anchoring axioms of the grammar enforce that positively and negatively marked nodes are paired off and the pairs identified. In the case of (4.46), the only way to do this that is in accordance with the tree axioms is given in (4.47).

(4.47) $r = 16 \wedge 11 = 17 \wedge 10 = 12 \wedge 13 = 14 \wedge 15 = 18 \wedge 19 = 20$

If our hearer adds this inferred information to his description of the discourse, the picture in (4.48) results.

[30] In Van Leusen and Muskens (2003) the root in a discourse description is anchored to a startsymbol, whose characterising property is the left closure statement. Furthermore, there is a distinction between T-nodes (for 'text') and S-nodes (for sentences or clauses), which is not maintained here.

(4.48) Δ_1^{rcl}:

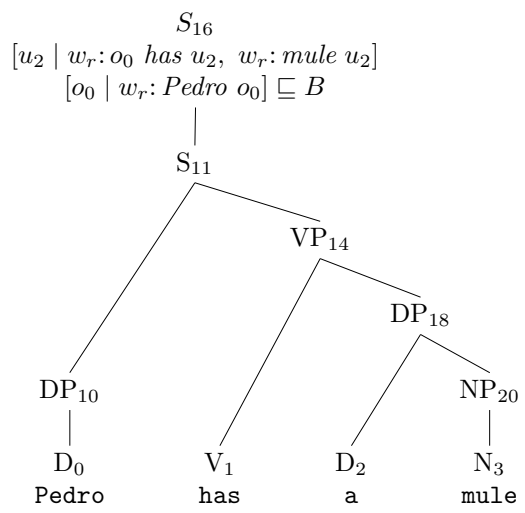

Note that in (4.48) not only the positively and negatively marked nodes have been 'clicked together' but that the identifications also allowed some semantic computation. The semantics of r can be computed to be $[u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2]$ and there still is a constraint $[o_0 \mid w_r\colon \mathit{Pedro}\ o_0] \sqsubseteq B$ on the background B.[31] These two pieces of information are typical of the kind of information that will be collected during the processing of any discourse. The semantics of the root, σ_r, will be interpreted as an *update* of the background B, while other information *constrains* the background. As we saw in section 4.2, the background must in a sense be constructed by the discourse participants. It is only partially constrained by what is said during a conversation and it will therefore remain largely underspecified. The meaning of the discourse given a certain background B will be identified with $B \oplus \sigma_r$, what was presupposed updated with what was said. In the present example this discourse meaning is $B \oplus [u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2]$.

Our hearer now has inferred certain information from the input description of sentence (4.44). But of course, the discourse may continue, in which case he will drop the right-closure hypothesis and increment the discourse description Δ_1 with an input description of the newly processed sentence. For example, suppose the discourse is continued with

(4.49) *He feeds it.*

After processing (4.49), the hearer may increment Δ_1 with the appropriate input description. In that case the following picture emerges.

[31] Remember the values of u are associated with *new* referents, while the values of o are referents that belong to the background B of the discourse. As will be explained in section 4.4.7, o-generated referents receive a referential interpretation in our semantics.

(4.50) Δ_2:

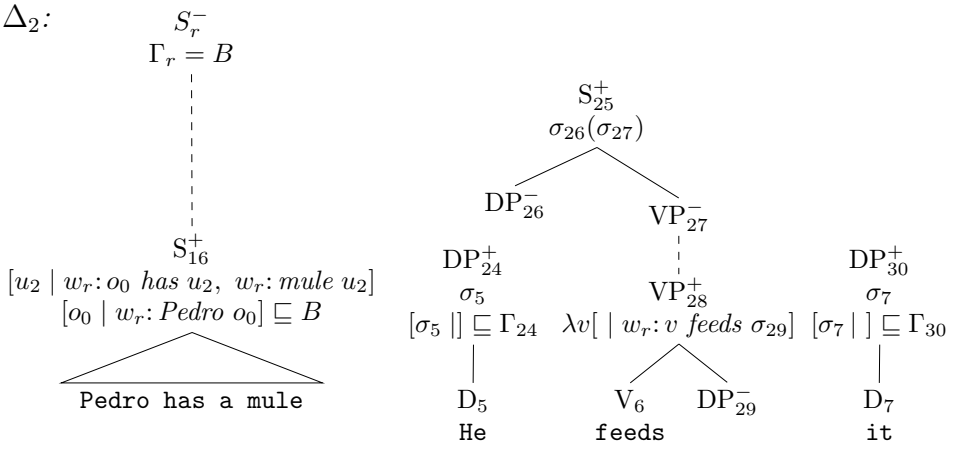

Note that the semantics of pronouns is not specified. While the indefinite a is associated with a discourse referent u_2, the constraint on the pronouns is that $[\sigma_k^\pi \mid] \sqsubseteq \Gamma_k$, where k is the terminal anchor of the elementary tree structure of the pronoun. This statement requires their semantics (a register) to be present in the universe of their local context.

Adding a new hypothetical end-statement to the discourse description, the hearer may reason about the syntactic and semantic composition of the new sentence and the discourse as a whole. As in our discussion in the preceding chapter, positively and negatively marked nodes in a discourse description of this kind can only be paired off in accordance with the axioms of the grammar if some implicit anchor introducing additional syntactic structure is present. Furthermore, the grammar allows only discourse relations to select implicit anchors. Thus, our hearer infers there is an implicit discourse relation which links the newly processed sentence to its discourse context through the elementary tree structure it anchors. On the basis of his world knowledge $\mathcal{K}$ and preferences induced by ■ he may select *elaboration* as the discourse relation most suitable to express the link of coherence between the new sentence and the preceding discourse.

(4.51) Δ_2^{rcl}:

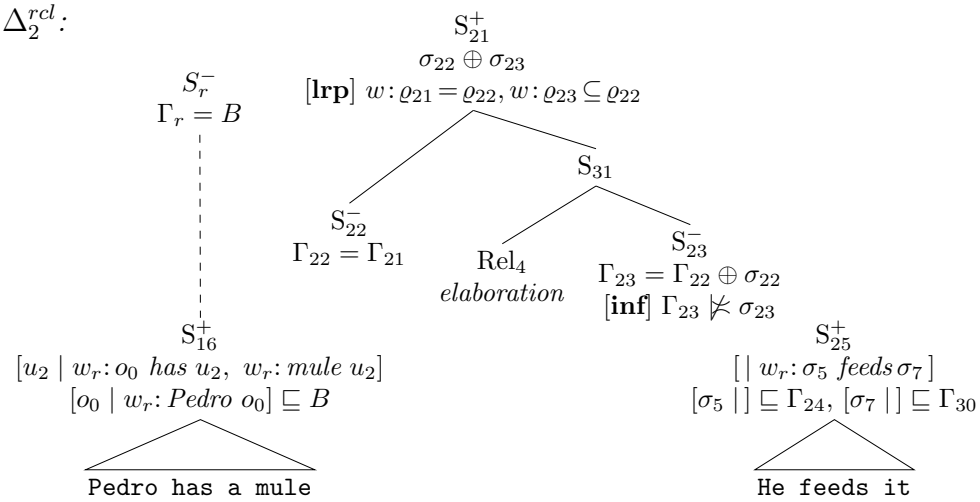

For expository purposes the discourse relation *elaboration* is temporarily enriched with a local informativity condition $\Gamma_{23} \not\models \sigma_{23}$. Again there is a unique way for the + and − nodes to click together and the picture below results. Here, with the help of $\mathcal{A}18$ and lexical information about Γ, the local context of each of the pronouns has been computed to be $B \oplus [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2]$, the discourse meaning of the previous chunk of discourse.

(4.52) Δ_2^{rcl}:

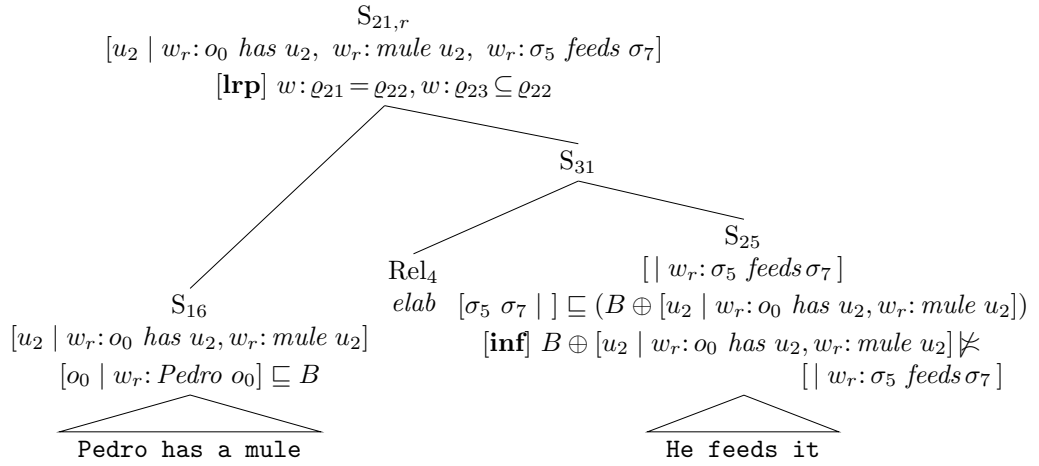

The constraints on spatio-temporal and mereological structure of discourse indicated by [**lrp**] are in fact part of σ_r and Γ_{25} (see section 4.4.4). Given an appropriate treatment in an eventuality semantic variant of the CDRT sublanguage, they should imply, in this particular discourse, that Pedro's ownership of a mule in part consists in feeding it. We will suppress this aspect of the interpretation in what follows. The discourse meaning which can be obtained from Δ_2^{rcl} is $B \oplus [u_2 |$ $w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7]$. As before, the meaning of the discourse is what was implicitly taken for granted (B), updated with what was said (σ_r). Both what was said and what was taken for granted, however, is constrained through the conditions generated by the various context sensitive elements in the discourse. Moreover, Global Consistency and Global Informativity (introduced in section 4.4.5) must be satisfied. The next section contains further discussion of these semantic constraints.

4.4.7 Testing Context Sensitive Constraints

Above we have seen how a hearer's reasoning from a discourse description resulted in the piecewise compositional construction of parts of a DRS. Here we shall see how further reasoning about the constraints on local contexts of various context sensitive elements, and general requirements on consistent and informative updating, results in a full or maximally specific discourse meaning given the discourse description. We will discuss the context sensitive constraints in turn, wrapping

up at the end of this section. As a preliminary to the discussion, we briefly zoom in on some properties of the implicit background B and the interpretation of discourse referents in the universe of the background.

Some Properties of the Background The background B will be constrained by context sensitive constraints that arise in the discourse. There are also a few general constraints, however, and these are summed up in $\mathcal{A}19$. We use $\mathsf{fst}(B)$ to denote the first element of the denotation of B, i.e. its universe; $\mathsf{snd}(B)$ would denote its second element, that is, its set of conditions.

$\mathcal{A}19$ $\mathit{proper}(B) \wedge \neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$

First, we take it that the background is *proper*. Modulo the naming of discourse referents, it should be possible for a background to have arisen out of previous discourse (with original background $[w_r \mid]$). Since discourse generates only proper DRSs, that is, DRSs in which each discourse referent occurring in a condition was declared in some universe accessible to it, backgrounds too should be proper. A second constraint is that the universe of the background does not contain registers that are values of u, as the latter will be reserved for discourse referents that arise during conversation. A last requirement on the background is that its universe contains w_r as an element.

Discourse referents that are present in the universe of the background arguably should be interpreted *referentially* rather than *existentially*. The referent w_r, which stands for the world of evaluation, may serve as an illustration, for *Pedro has a mule* should not be interpreted as 'there is some world in which Pedro has a mule' but as 'Pedro has a mule in *this* world.' The discourse referent corresponding to *Pedro* should also be interpreted as given and similarly should deictic pronouns and perhaps even definite descriptions whose referent is accommodated.

A notion of truth that treats all background referents referentially will result if the following definition of truth is adopted:[32] a DRS K is said to be *true in state i with respect to background B* if $\mathit{true}(\langle \mathsf{fst}(K) - \mathsf{fst}(B), \mathsf{snd}(K)\rangle)(i)$ holds, that is, referents present in the background's universe will not be interpreted existentially but will receive their interpretation from the state i. It will be assumed that discourse participants interpret the discourse with respect to some state i_0 that they take the discourse to be *about*[33] and we will often be interested in the truth of some DRS K in i_0 with respect to a given background.

[32] See appendix A for the original truth definition in fine-grained CDRT, of which this is a revision.

[33] Austin (1961): "A statement is said to be true when the historic state of affairs to which it is correlated by the demonstrative conventions (the one to which it 'refers') is of a type with which the sentence used in making it is correlated by the descriptive conventions." In the present context, i_0 (or its sequence of values) can be compared to Austin's 'historic state of affairs'. Presumably, our states are somewhat richer than Austin's states of affairs, for besides a possible world they contain values for all kinds of referents. On Austinian propositions, see also (Barwise and Perry 1983).

To illustrate this further, let us consider the DRS $K =$

$$[w_r\, o_1\, u_3 \mid w_r{:}\, \mathit{Pedro}\ o_1,\ w_r{:}\, o_1\ \mathit{has}\ u_3,\ w_r{:}\, \mathit{mule}\ u_3]$$

and assume that w_r and o_1, but not u_3, are members of the universe of background B. If, in general, we agree, for readability, to write δ^0 for $V(\delta, i_0)$, we can express the truth of K in i_0 with respect to B as follows. The function **wp** translates a formula in the CDRT sublanguage into predicate logic, see appendix A.

$$\begin{aligned}\mathbf{wp}([u_3 \mid w_r{:}\, \mathit{Pedro}\ o_1,\ w_r{:}\, o_1\ \mathit{has}\ u_3,\ w_r{:}\, \mathit{mule}\ u_3], \top)^{i_0} = \\ \exists x_3\, [\mathit{Pedro}\ (o_1^0, w_r^0) \wedge \mathit{has}\ (o_1^0, x_3, w_r^0) \wedge \mathit{mule}\ (x_3, w_r^0)]\end{aligned}$$

Here the values of o_1^0 and w_r^0 depend on the state i_0, which functions much as an *external anchor* in the sense of Kamp and Reyle (1993), pp. 246–248.[34]

With the notion of background now elucidated, let us see how reasoning from input descriptions can result in a process familiar from DRT: While indefinites create fresh referents, pronouns will pick up old referents (or must get a deictic interpretation). The material connected with proper names is relegated to the main Discourse Representation Structure. Constraints on local contexts must either be satisfied or accommodated. Global Consistency and Informativity when satisfied constrain global accommodation.

Proper Names Proper names end up in the main DRS. The lexical description for *Pedro* contains the statement in (4.53a), requiring the semantic material connected with the name to be in the background.

(4.53) a. $[o_0 \mid w_r{:}\, \mathit{Pedro}\ o_0] \sqsubseteq B$

b. $[w_r\, o_0 \mid w_r{:}\, \mathit{Pedro}\ o_0] \sqsubseteq B$

c. $\sigma_r = [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7]$

d. $[w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7] \sqsubseteq B \oplus \sigma_r$

In fact, using $\mathcal{A}19$, this can be strengthened slightly to (4.53b). Using (4.53c) (which was found in (4.52)), the hearer can now derive (4.53d). The discourse meaning $B \oplus \sigma_r$ must contain (and therefore entail) a certain DRS containing the material connected with the name. It should be clear that this is in fact an implementation of global accommodation.

[34] While Kamp and Reyle (1993) use anchors only for proper names, our use is wider as we let i_0 interpret all referents in the discourse's background. There is some room for fine-tuning here and it may also be held that some referents from the background are referential while others are interpreted existentially. See section 6.2.4 in chapter 6.

Pronouns Indefinites create fresh referents, but pronouns pick up old referents. That indefinites create fresh referents is a property they share with many other expressions. Their creation is a by-product of the perpetual creation of fresh node names. The referent u_2 connected with the indefinite *a* in (4.52) was created in this way. The moment that it was established that there was some node 2 carrying the lexeme *a* the referent u_2 also sprung into existence. Note that, by the injectivity of u, u_2 cannot corefer with a referent created at any other node.[35] Since, by $\mathcal{A}19$, u_2 can not be an element of the universe of B and since, by the same axiom, B is proper, there can also be no condition in B in which u_2 occurs. The referent is therefore truly fresh to the discourse.

The process whereby pronouns get bound by referents that already exist takes a little more care to explain. In (4.54a) a constraint is shown that was collected in (4.52). It resulted from the requirement on pronouns that $[\sigma_k^\pi \mid] \sqsubseteq \Gamma_k$ (with k the node carrying the pronoun). This was required of the *he* and *it* nodes and in both cases the local context Γ_k could be computed to be $B \oplus [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2]$. Using (4.53a) and the idempotency of $\oplus$ it is seen that in fact (4.54b) must hold.

(4.54) a. $[\sigma_5\ \sigma_7 \mid] \sqsubseteq B \oplus [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2]$

b. $[\sigma_5\ \sigma_7 \mid] \sqsubseteq B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2]$

This puts constraints on what σ_5 and σ_7 are, but there are still many possibilities for these referents to resolve. In principle this is what we want, as linguistic information typically underspecifies how pronouns should resolve. However, the constraints collected thus far leave open too many possibilities and therefore more constraints are needed.

In order to illustrate this, let us concentrate on σ_7 and see what values it can take. One possibility is that the pronoun *it* has no linguistic antecedent and that σ_7 is to be identified with some register that was in the universe of the background B but is not mentioned in the discourse. We think this possibility is in fact welcome and corresponds to a reading of the text in which *it* refers to an object that is somehow salient (for example, as a result of pointing) but that was not introduced by linguistic means. We identify this with the deictic use of the pronoun.

Technically, there is also the possibility that σ_7 is in fact equal to w_r. This possibility is an artefact and a consequence of our choice to have one type of register for worlds and individuals. We exclude it by postulating that no node can have w_r as its semantics.

$\mathcal{A}20\ \neg\exists k\ \sigma_k^\pi = w_r$

[35] Of course the *value* of another referent may be identical with the value $V(u_2, i)$ of u_2 at any given state i. The difference between identity of referents and identity of their values is crucial to the present discussion.

A fuller treatment should perhaps comprise a more general type or kind distinction of the registers involved, so that $\mathcal{A}20$ would fall out as a consequence.

Of the remaining possibilities, the identification $\sigma_7 = o_0$ should be excluded on linguistically more interesting grounds. Since *Pedro* has masculine gender, it cannot antecede *it*, which is neuter. Since this does not follow from the theory we have set up thus far, and since it is of course just one manifestation of a phenomenon that really pervades language, we must add a general mechanism for feature constraints. In fact this is very easy and Muskens (2001) shows how it can be done on the basis of the first-order feature theory of Johnson (1991). We refer to the discussion in Muskens (2001) for details. Using the notation of that paper, a general requirement to the effect that coreferring nodes should agree (on number and gender) could be formulated as follows.

$$(4.55)\quad \forall k k' \, [\sigma_k^\pi = \sigma_{k'}^\pi \rightarrow \forall f \, [\mathrm{arc}(\boxed{k}, \textsc{agr}, f) \rightarrow \mathrm{arc}(\boxed{k'}, \textsc{agr}, f)]]$$

This could be added as an axiom, together with the three axioms (axiom schemes) for features in Muskens (2001). Note the mixed semantic / syntactic character of (4.55). This is a case of genuine mutual constraint between syntax and semantics.

With a mechanism for features such as that in place, a general requirement to the effect that coreferring nodes should agree (on number and gender) could be added as a syntactic-semantic axiom to the grammar. Given the implementation of this requirement, there are two possibilities left for the denotation of σ_7: $\sigma_7 = u_2$ or $\sigma_7 \notin \{w_r, o_0, u_2\}$. Similarly the hearer can deduce that either $\sigma_5 = o_0$ or $\sigma_5 \notin \{w_r, o_0, u_2\}$. In fact the following four possibilities remain.

(4.56)a. $[w_r\, o_0\, u_2 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon o_0\ \mathit{feeds}\ u_2]$
$\sqsubseteq B \oplus \sigma_r$

b. $[w_r\, o_0\, u_2\, \sigma_7 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon o_0\ \mathit{feeds}\ \sigma_7]$
$\sqsubseteq B \oplus \sigma_r \wedge [\sigma_7 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ]$

c. $[w_r\, o_0\, u_2\, \sigma_5 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon \sigma_5\ \mathit{feeds}\ u_2]$
$\sqsubseteq B \oplus \sigma_r \wedge [\sigma_5 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ]$

d. $[w_r\, o_0\, u_2\, \sigma_5\, \sigma_7 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon \sigma_5\ \mathit{feeds}\ \sigma_7]$
$\sqsubseteq B \oplus \sigma_r \wedge [\sigma_5 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ] \wedge [\sigma_7 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ]$

The first possibility corresponds to the preferred reading (Pedro feeds the donkey), the second to the case where *it* was taken deictically, the third to a deictic reading of *he*, and the fourth to the case where both pronouns pick up an extralinguistic referent. The hearer must now make a choice as to what was meant on the basis of his world knowledge and of what he knows or may assume is in the background. Note that each four of the possibilities offers suitable input for a translation into predicate logic,[36] and that in each case certain conclusions can be drawn from

[36]See Definition A.4 in appendix A. Furthermore, in each case all relevant discourse markers can be shown to be disequal (the precondition to Theorem A.5).

the assumption that $B \oplus \sigma_r$ is true in i_0 with respect to background B. These conclusions are, respectively

(4.57) a. $\exists x\,[\mathit{Pedro}\ (o_0^0, w_r^0) \wedge \mathit{has}\ (o_0^0, x, w_r^0) \wedge \mathit{mule}\ (x, w_r^0) \wedge \mathit{feeds}\ (o_0^0, x, w_r^0)]$,
b. $\exists x\,[\mathit{Pedro}\ (o_0^0, w_r^0) \wedge \mathit{has}\ (o_0^0, x, w_r^0) \wedge \mathit{mule}\ (x, w_r^0) \wedge \mathit{feeds}\ (o_0^0, \sigma_8^0, w_r^0)]$,
c. $\exists x\,[\mathit{Pedro}\ (o_0^0, w_r^0) \wedge \mathit{has}\ (o_0^0, x, w_r^0) \wedge \mathit{mule}\ (x, w_r^0) \wedge \mathit{feeds}\ (\sigma_5^0, x, w_r^0)]$,
d. $\exists x\,[\mathit{Pedro}\ (o_0^0, w_r^0) \wedge \mathit{has}\ (o_0^0, x, w_r^0) \wedge \mathit{mule}\ (x, w_r^0) \wedge \mathit{feeds}\ (\sigma_5^0, \sigma_7^0, w_r^0)]$.

By evaluating the discourse meaning $B \oplus \sigma_r$ relative to his world knowledge the hearer may generate preferences over the different possible readings, in other words further narrowing down of the possibilities may be a result of pragmatic reasoning. Deictic readings will only be preferred if some salient entity is present in the background.

While in this example only linguistically acceptable pronoun resolutions were left, other suitability constraints are necessary in other cases. One set of constraints that immediately springs to mind is that of the syntactic Binding Theory (Chomsky 1981). As we shall see in chapter 5 the principles of Binding Theory, and the geometric notions in terms of which it is stated, for example c-command, can straightforwardly be formulated in LDG. The description approach allows us to integrate theories from various levels and a syntactic theory such as the Binding Theory, if adopted, has immediate relevance for the construction of Discourse Representation Structures. A related issue is the interaction with discourse syntactic structure, or rather, discourse topical structure as induced by discourse syntax. We will touch upon this in chapter 5.

Appropriateness Conditions on Local Contexts Appropriateness conditions attached to discourse relations or connectives must be satisfied; the effect of this is that the implicit background B is constrained. We let the discourse relation *elaboration* in (4.52) introduce an informativity constraint on the local context of the discourse unit it links to the discourse context. The constraint is given in (4.58a). It says that certain descriptive material, corresponding to the semantic value of the relevant discourse unit, must not be entailed by the relevant local context. Thus it enforces that this material is *new* or informative in the discourse.

(4.58) a. $B \oplus [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2] \not\models [\,|\, w_r{:}\, \sigma_5\ \mathit{feeds}\, \sigma_7\,]$
b. $B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2] \not\models [\,|\, w_r{:}\, \sigma_5\ \mathit{feeds}\, \sigma_7\,]$
c. $B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2] \not\models [\,|\, w_r{:}\, o_0\ \mathit{feeds}\, u_2\,]$

Using (4.53a), $\mathcal{A}19$, and the idempotency of $\oplus$ the requirement can be strengthened to (4.58b). Obviously, the requirement must be satisfied given some resolution of the pronouns in 'he feeds it'. Though in the example at hand, all of the

possible readings summed up in (4.56) allow for the satisfaction of (4.58b), in principle, resolutions of the pronouns that would prevent its satisfaction would be excluded. They would result in uninformative discourse. (4.58c) represents the informativity condition given the resolution of the pronouns in (4.56a), the most preferred reading. Since it holds that

$$[w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2] \not\models [\,\mid w_r{:}\, o_0\ \mathit{feeds}\, u_2\,]$$

satisfaction of (4.58c) implies that the implicit background B contains no information such that $[\,\mid w_r{:}\, o_0\ \mathit{feeds}\, u_2\,]$ can be inferred. For example, B cannot contain the information that if Pedro has a mule, he feeds it, or that mules are animals and Pedro feeds all animals. The local informativity condition prevents the accommodation of such material in B. In this way it constrains the discourse meaning $B \oplus \sigma_r$.

Global Informativity Global Informativity defines informativity at the level of discourse contributions, or updates of discourse descriptions. It constrains input discourse meanings. In the example at hand, Global Informativity results in the following constraint.

There must be a DRS K such that
$\Delta_1^{rcl} \not\models (B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0, w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2] \mid\!\sim K)$ and
$\Delta_2^{rcl} \models (B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0, w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2\ w_r{:}\, o_0\ \mathit{feeds}\, u_2] \mid\!\sim K)$

This can be satisfied in two situations. Either, there is a DRS K such that $\Delta_1^{rcl} \not\models (K \sqsubseteq B)$ and $\Delta_2^{rcl} \models (K \sqsubseteq B)$. This means that due to processing 'he feeds it' there is information K which is accommodated in B but which is not yet entailed or accommodated in B in Δ_1. Global Informativity is satisfied because of the information growth induced by accommodation. The other possibility is that part or whole of the semantic content of the contribution is new relative to the input discourse meaning. Suppose $K = [\,\mid w_r{:}\, o_0\ \mathit{feeds}\, u_2]$. It is clear that for this choice, the second clause of the Global Informativity requirement is satisfied. The first clause says that

(4.59) $\Delta_1^{rcl} \not\models$
$B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2] \mid\!\sim$
$[\,\mid w_r{:}\, o_0\ \mathit{feeds}\, u_2]$

This prevents the accommodation in B of information that would support the inference of $[\,\mid w_r{:}\, o_0\ \mathit{feeds}\, u_2]$, as with the local informativity constraint discussed earlier. It seems reasonable to assume that if no additional constraints, specified in Δ_2, change the hearer's assumptions about B, the constraint in (4.59) carries over to Δ_2.

Starting out from the informativity condition [**inf**] in Δ_2 we can reason as follows. Either it induces accommodation in B or it doesn't. If it induces accomodation, Global Informativity is satisfied. If it does not induce accommodation, it follows that the constraint must already be satisfied in Δ_1, that is, we have

(4.60) $\Delta_1^{rcl} \models$
$B \oplus [w_r\, o_0\, u_2 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2] \not\models$
$[\ \mid w_r\colon o_0\ \mathit{feeds}\, u_2]$

The first clause of the Global Informativity requirement (4.59) follows from this. The second clause also holds, so Global Informativity is satisfied. In sum, if we choose K such that it is identical to the semantic content of the contribution, then if the local informativity condition attached to *elaboration* is satisfied, Global Informativity is satisfied as well. In the case at hand either the one or the other requirement is superfluous. As discussed in section 4.4.5, however, it is a matter for further investigation whether in general, both types of informativity requirements should be maintained, or one may be dropped in favour of the other.

Global Consistency We have seen that all context-sensitve constraints triggered by the lexemes and implicit elements in the discourse under consideration can be satisfied, resulting in a partial specification of the implicit background and in four possible readings of the discourse. Evaluating the discourse meaning relative to his world knowledge, the hearer may select from these a single most preferred reading. For this reading, as for all grammatical readings, Global Consistency requires that the discourse meaning be consistent. Thus, in addition to the following condition must be satisified.

(4.61) $B \oplus [w_r\, o_0\, u_2 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\, u_2,\ w_r\colon \mathit{mule}\ u_2\ w_r\colon o_0\ \mathit{feeds}\, u_2] \not\models \bot$

As with the informativity condition attached to the discourse relation *elaboration* and the Global Informativity condition, (4.61) constrains the implicit background by preventing the global accommodation of any material that would result in an inconsistent discourse meaning.

Concluding, it was shown that discourse meanings are constrained in various ways by context sensitive elements occurring in the discourse, and by global semantic conditions. Names trigger conditions on the implicit background B so that they either resolve to, or are accommodated in the main DRS. Pronouns introduce conditions on their local context which force their semantic value to either unify with an existing discourse referent, or be globally accommodated and get a deictic interpretation. Appropriateness conditions triggered by discourse relations must be satisfied in the relevant local context. As a consequence, they constrain the implicit background. Global Informativity constrains the relation between input,

and output discourse meanings. Its predictions overlap with those of local informativity constraints. Global Consistency, by requiring the (output) discourse meaning to be consistent, constrains the implicit background as well.

All of these constraints contribute to the specification of (underspecified) discourse meanings. Since language abounds with context sensitive elements, in a fully fledged discourse grammar there will be many more lexical elements that come with additional constraints on local contexts. Notably, a lot of lexemes introduce *presuppositions* of one kind or another. In fact, the appropriateness conditions attached to discourse relations and discourse operators discussed in section 4.4.4 may be viewed as linguistic presuppositions. In chapter 6 we shall zoom in on presuppositions attached to definite descriptions, and discuss a treatment that is in some respects similar to that of pronouns and in other respects to that of appropriateness conditions. In all these cases, context sensitive aspects of meaning are captured in the lexical semantics. By integrating local contexts in the lexical descriptions of these elements, a straightforward declarative treatment of context sensitivity is obtained.

It remains to be seen how much of context sensitivity in discourse interpretation can or should be accounted for in the lexicon of the grammar. Consider the Global Consistency condition, one of the two nonlexical constraints discussed above. Though the effect of this condition could also be obtained by including consistency requirements on local contexts in the lexical descriptions of discourse relations and discourse operators, we would lose a generalisation in doing so. This may be the case for other felicity conditions attached to discourse contributions as well. However, for any form of context dependence of interpretation normally captured under the heading of anaphoricity or linguistic presupposition we expect that it can be treated in the lexicon of the grammar. We expect the same for the context sensitive constraints triggered by the information structure or topic-focus structure of clauses.

All in all, local contexts and conditions on them can simply be treated as an additional parameter in the semantic composition of discourse; the result may thus be called a context sensitive compositional discourse semantics.

4.5 Conclusion

In the preceding chapters we built a framework of discourse interpretation around a logical description grammar or LDG. For an interpreting agent or language user a, a framework of discourse interpretation $\mathcal{F}_a = \langle G^+, \mathcal{K}, \blacksquare \rangle$ consists of an incremental LDG system for discourse analysis, a reservoir of world knowledge, and a default reasoning module. The grammar system G^+ combines a discourse grammar $\mathcal{G}$ with discourse descriptions that are built up incrementally from input or sentence descriptions. $\mathcal{G}$ represents the language user's grammatical knowledge of a given natural language, $\mathcal{K}$ his nonlinguistic, general knowledge or world knowl-

edge, and ■ stands for his default reasoning capacity and his capacity to assign preferences over verifying tree models of discourse descriptions. Thus, $\mathcal{F}_a$ as a whole is a knowledge system.

The framework distinguishes two levels of linguistic analysis: the level of possibly underspecified discourse descriptions and the level of the annotated tree structures or models that 'fit' or verify these descriptions. For any given stretch of discourse a discourse description is generated. This is a set of statements in type logic, describing the syntactic and semantic relations characterising the discourse processed so far. Discourse descriptions figure as 'underspecified representations' of the discourse. They represent the language user's knowledge of the ongoing discourse. The discourse grammar itself is also a set of logical statements in type logic. A language user may reason about the models or tree structures that verify a given discourse description given the grammar. The tree structures feature as fully specified representations of the discourse processed so far. Since a discourse description can have more than one verifying tree model, the syntax and semantics of the discourse may remain underspecified.

Some of this underspecification may be of the 'don't care' kind: the language user may be able to draw all conclusions that he desires without further filling in the details. But in many cases further specification is necessary and should come from pragmatic considerations, general felicity conditions on discourse interpretation, interpretational constraints that are contributed by intonation, and so on. Pragmatic considerations are covered in our framework to the extent that a language user may evaluate a given (underspecified) discourse meaning relative to his world knowledge $\mathcal{K}$ and generate preferences over different possible readings employing his default reasoning engine ■. In how far aspects of interpretation commonly categorised as 'pragmatics'—such as the inference of implicatures, or the conditions of appropriate use attached to discourse contributions—can or should be captured in a preference system remains to be seen, however. Constraints covered by intonation are not covered in the grammar as yet. Though we do not expect special problems in letting discourse descriptions talk about intonation, clearly there is still an area to be explored here.

Our focus of attention is on discourse syntax and semantics. The discussion in this chapter centered on discourse meaning and the integration of context sensitive constraints in the grammar. The meaning of a discourse is defined as $B \oplus \sigma_r$, the update of an underspecified implicit background B with σ_r. This background represents what is taken for granted by the participants, as constrained by what is said in the conversation. Discourse meanings function as context of interpretation for the next discourse contribution. General conditions on consistent and informative updating were formulated as constraints on discourse meanings and added as semantic axioms to the grammar. Much of the language user's reasoning in obtaining a discourse meaning, however, concerns the interpretation of context sensitive elements in a discourse contribution, such as anaphors, presupposition triggers, discourse relations and discourse operators. The context sensitivity of

these elements is captured by integrating local contexts, and constraints on them, in their lexical descriptions. The description method results in a unified, declarative treatment of context sensitivity.

The descriptions perspective on syntax and semantics advocated in these chapters leads to a considerable change in the overall logical architecture of linguistic theory. Nevertheless it remains compatible with most existing ideas in mainstream linguistics and can even be argued to make for a smoother integration of those existing ideas. The overall logical architecture is changed because the theory focuses on descriptions of linguistic objects (including meanings and the form-meaning relation) rather than on those linguistic objects themselves; it also describes the syntactic and semantic levels of the grammar in parallel and does not follow the more common pipelined architecture in which syntactic forms must be produced before interpretation can take place. Thus room is created to account for an interdependency of syntax and semantics. On the other hand, the view that language users represent language by means of descriptions rather than by means of structures does not necessarily lead to any drastic revision of our understanding of what these structures are. In this chapter combinations of standard Discourse Representation Structures and standard phrase structure trees were taken to be the objects of description. The theory obviously is also compatible with many other choices, but no radical departure from common practice is needed or indeed desired in this respect.

In the preceding chapters the LDG framework of discourse interpretation was presented as a work-bench for linguists, a logical tool for the specification of linguistic theories. The strength of the formalism, we believe, is in its treatment of the *interaction* of the various linguistic and nonlinguistic parameters that govern the interpretation of discourse, and in the fact that underspecification of linguistic information is a natural consequence of the approach. Underspecification or implicitness in discourse is all-pervasive. Nonetheless, language users generally arrive at sufficiently specified readings of discourse, seemingly without great effort. It is the combination of these facts that must be explained in a discourse theory. The LDG framework, with its emphasis on the embedding of grammatical knowledge within the language user's general knowledge, provides a platform for this.

By way of testing the framework, in what follows it will be applied to presupposition theory. As a preliminary, we will zoom in on anaphora resolution in chapter 5. We put forward an analytic variant of the treatment of pronoun resolution that was offered in this chapter, the advantage of which is that it models anaphoricity more clearly as an independent parameter in the discourse theory. In terms of the predictions of the theory, there are no differences, however. Furthermore, the interaction with syntactic and semantic constraints, salience and topicality effects, and accommodation will be more fully addressed.

Chapter 6 then provides an implementation of presupposition theory in LDG. Our aim is not so much to spell out a new treatment of presupposition, but to

show that central aspects of existing theories can fruitfully be combined, once they are embedded in a sufficiently powerful discourse model, which incorporates a sufficiently fine-tuned discourse semantics. The approaches that will be combined are the one based on Karttunen (1974), which models presuppositions as tests on contexts of interpretation, and the anaphoric approach to presupposition projection put forward in Van der Sandt (1992), which binds or accommodates presuppositional material in accessible locations in discourse meanings. It will be shown that, given a certain choice of preferences in the framework of discourse interpretation, van der Sandtian predictions may in principle be obtained, while the conceptually attractive analysis of presuppositions as tests on the context of interpretation is maintained. The interaction with independent constraints in the discourse grammar, and with the hearer's reasoning about world knowledge and his preferences is crucial for the explanation.

Chapter 5

Anaphoricity

5.1 Introduction

A main source of coherence in discourse is anaphoricity. *Anaphors* or *anaphoric expressions* are lexical elements in natural language whose common characteristic is that they denote entities or properties which were mentioned earlier in the discourse, or are present in the nonlinguistic context, or can be inferred or assumed to be present in the context of interpretation on the basis of the linguistic and nonlinguistic context of interpretation. The term *antecedent* refers to a linguistic expression in the discourse which, if present, usually precedes the anaphor and which is responsible for introducing the entity it denotes in the discourse context. Where anaphors that denote individual entities are concerned, we use the term *binder* to indicate a discourse marker declared in the universe of the context of interpretation whose occupant (or the value assigned to it) is the individual entity the anaphor denotes.

Finding out the meaning of a newly uttered sentence involves resolving the anaphors occurring in it, relative to the discourse context established by the preceding discourse. In the context of LDG, *resolving* an anaphor means determining its potential binders and selecting the binding relations that result in most preferred readings of the discourse. Nominal anaphors introduce underspecified discourse markers which must be bound by a semantically accessible binder in the discourse meaning. We will propose that binding is the effect of general semantic conditions on the properness of discourse meaning and local contexts (section 5.2). To account for context-sensitive elements whose binder is introduced via subsequent discourse (often called 'kataphors'), a form of lazy interpretation may be assumed. Furthermore, predicational anaphors, such as VP-ellipsis and so-called 'pay-check pronouns' may resolve as the result of parallelism requirements rather than properness.

Differences between various classes of anaphors are laid down in the lexicon.

Specific syntactic, semantic and informational conditions on its relation to the discourse context may characterise a given type of anaphor. To illustrate the classifying effect of additional syntactic constraints, we briefly discuss an implementation of the Chomskian Binding Conditions in LDG (section 5.3). Furthermore, we introduce the notion of 'discourse availability' (section 5.4) to get a grasp on the constraining effect of discourse topical structure and salience effects on anaphora resolution. One of the features that distinguishes personal pronouns from, for example, definite descriptions is the fact that they require binders that are *most salient* given the discourse topic hierarchy. Definite descriptions lack this requirement, though they seem to be sensitive to the topical status of their binders in some measure.

Classifying conditions of anaphoric elements not only explain what sort of discourse context licences the use of what kind of anaphor, they also account for differences in the anaphor's potential to resolve to implicit background information and induce the accommodation of a binder. Since Lewis (1979), the term *accommodation* is employed in semantic theory to denote the implicit adjustment of the context of interpretation, as induced by the interpretation of a statement made relative to that context. In the LDG formalism it can be modeled as exactly that: accommodation is the specification of hitherto underspecified background information, evoked by the interpretation of the ongoing discourse.[1] Anaphors and presuppositional expressions may resolve relative to what has been explicitly stated in the discourse, but they may also teach the hearer something new about what is implicitly presupposed. We will zoom in on the accommodation potential of pronouns, and propose that the salience requirement on their binder is what makes accommodation of a binder costly. Therefore it is usually dispreferred (section 5.5).

In all of this, what makes a hearer resolve anaphoric links is his ability to *reason* about what is stated in the discourse given his common knowledge. While binding relations as such receive a purely declarative treatment in our grammar, anaphora resolution is essentially an inference process, and that is how we model it. Some of the information on the basis of which the hearer reasons is provided linguistically, other information is part of his common knowledge. Given the architecture of our framework of discourse interpretation the language user may access these information sources in parallel.

[1] As in the preceding chapter, discourse meanings are defined as $B \oplus \sigma_r$, the implicit background updated with the semantic content of the discourse. For the time being, accommodation is nothing but the acceptance of information in B. This implements only *global* accommodation, that is, the specification of the initial context of interpretation in possible world approaches, or of the main DRS in DRT. In the second half of this chapter, a generalisation to local and intermediate accommodation will be provided.

5.2 Properness and Pronoun Resolution

The analysis of pronouns presented in the preceding chapter may be summed up as follows. Pronouns are context-sensitive elements. Their semantic values are underspecified discourse markers. They introduce a context-sensitive constraint which requires that the discourse marker which is their semantic value is present in the universe of their local context (see the lexical description in (5.1a)). Thus, pronouns must be bound in their local context. Resolving a pronoun comes down to finding out with which of the discourse markers present in the local context the underspecified value can be identified and selecting one that results in a most preferred reading.

(5.1) a. $\mathrm{DP}^{+}_{k_1}$, σ^{π}_{k}, $[\ \sigma^{\pi}_{k} |\ ] \sqsubseteq \Gamma_k$, D_k, he

b. $\mathrm{DP}^{+}_{k_1}$, σ^{π}_{k}, D_k, he

We shall follow a slightly different approach here, and propose a variant of our treatment which results in the same predictions, but which highlights that binding as such is an independent parameter in the discourse semantics. The fact that pronouns must be bound in their local context shall be made to follow from general requirements on properness, rather than from lexical constraints. Instead of descriptions such as (5.1a) we adopt descriptions of the kind illustrated in (5.1b), which only state that the semantic value of a pronoun is a discourse marker. Simultaneously, we add the following general axiom to the grammar. We shall often refer to this as the *Properness condition*.[2]

$\mathcal{A}21$ $\mathit{proper}(B \oplus \sigma_r) \wedge \forall k[\,\mathit{proper}(\Gamma_k)]$

This says that, for any discourse description, the discourse meaning $B \oplus \sigma_k$, and the local contexts of all nodes in the described structure must be proper DRSs. As in the preceding chapter, a DRS K is *proper* if it has no free referents, and the *free* referents of a DRS K are the ones that are free in some condition of K but are not in the universe of K. The effect of $\mathcal{A}18$ is that all discourse markers employed in some condition in $B \oplus \sigma_r$ or in arbitrary local contexts are declared in an accessible universe, in other words they are *bound*. Since every discourse marker employed in the semantic value of a discourse unit figures in the local

[2]The condition subsumes the first part of axiom $\mathcal{A}19$, introduced in the preceding chapter, which states some properties of the implicit background: $\mathit{proper}(B) \wedge \neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$. This may now be reformulated as

$\mathcal{A}19'$ $\neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$

.

context of the consecutive discourse unit, or, if the discourse unit is the currently rightmost one of the discourse, in the resulting discourse meaning, all pronouns whose semantic value ends up in the semantic value of the sentence that contains them are subject to axiom $\mathcal{A}21$. Hence, at every stage in processing a discourse, these pronouns are bound in their local context.[3]

The properness condition expresses the intuition that discourse meanings resulting from grammatical discourse, and contexts of interpretation in general, are well-formed in the sense that they correspond to closed formulas or sentences in predicate logic. In fact, it is a common assumption in DRT that output DRS's resulting from processing and interpreting a grammatical sentence, are proper.[4] The first half of our axiom reflects this. However, whereas in standard DRT properness is a metalevel property of DRS's, in LDG 'properness' is made availabe in the semantics at the level of the describing logic, see appendix A. It is an object level property which the language user's grammar can 'talk about'. As such, it may naturally be employed to specify or describe parts of his grammatical knowledge, notably presupposition theory.

We will in addition assume that properness is a built-in characteristic of the object level relation of entailment. In grammatical discourse we are only concerned with testing entailment relations $\Gamma \models\!\!\!\sim K$ where there are no discourse markers free in K. A slightly modified definition of object level entailment which takes care of this is provided in appendix A. Given this modification we can be sure that all discourse markers introduced by anaphoric elements, either via their semantic content or embedded in presuppositional material, are subject to a properness condition and must be bound in the local context of the anaphor. As before, the semantic values of pronouns are underspecified discourse markers and pronoun resolution is reasoning about potential binding relations, but now the language user infers that the semantic value is bound in its local context on the basis of his general grammatical knowledge rather than lexical constraints. Binding is an effect of a basic semantic well-formedness constraint on the discourse meaning and local contexts of interpretation. The account carries over to all lexical elements that introduce free discourse markers as part of their semantics.

5.2.1 Inferring Binding Relations

Let us see how this slightly modified treatment of pronoun resolution works out for the example that was discussed in section 4.4.6 of chapter 4 before. Consider

[3]Arguably, just requiring $\mathit{proper}(B \oplus \sigma_r)$ is enough. But then an integrity rule or preference should be added to the framework of discourse interpretation to prevent that bindings of discourse markers that are inferred at one stage in the interpretation of a discourse are changed after incrementation steps, when new antecedents have been made available through subsequent discourse.

[4]Compare (Kamp and Reyle 1993), section 1.4. Furthermore, Van der Sandt (1992) states the properness requirement on output DRSs as one of the constraints on anaphora binding.

the following discourse.

(5.2) *Pedro owns a mule. He feeds it.*

After processing the second sentence in (5.2), the hearer may increment the discourse description established so far with the appropriate input description. The resulting discourse description can be graphically represented as follows.

(5.3) Δ_2:

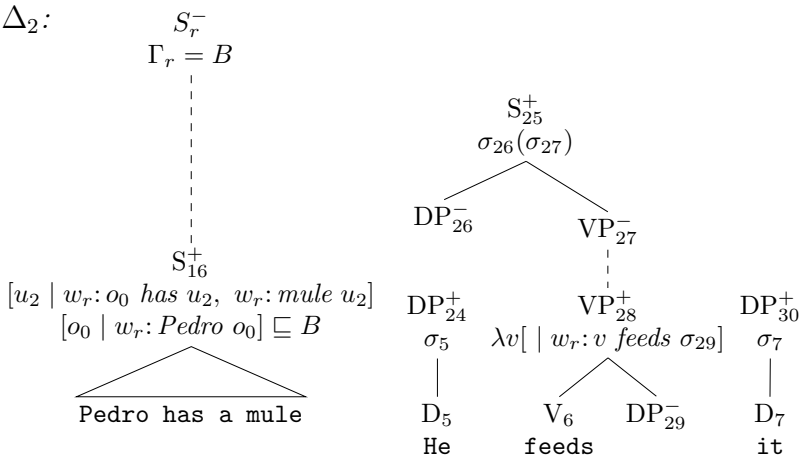

Notice that while the name 'Pedro' is connected with a discourse referent o_0 that is required to be present in the universe of the implicit background B, and the indefinite 'a mule' creates a fresh discourse referent u_2, the pronouns in 'he feeds it' introduce the underspecified discourse referents σ_5^π and σ_7^π. Their binding must be inferred from general axioms in the grammar.

Reasoning in the by now familiar way, the language user may pair off positively and negatively marked nodes, and infer the presence of an implicit discourse relation. Discourse relation *elaboration* is a plausible candidate in the example at hand. Furthermore, temporarily closing off the discourse description, he may compose the semantic values of the nodes. The following picture of the discourse description shows the result of this reasoning.

(5.4) Δ_2^{rcl}:

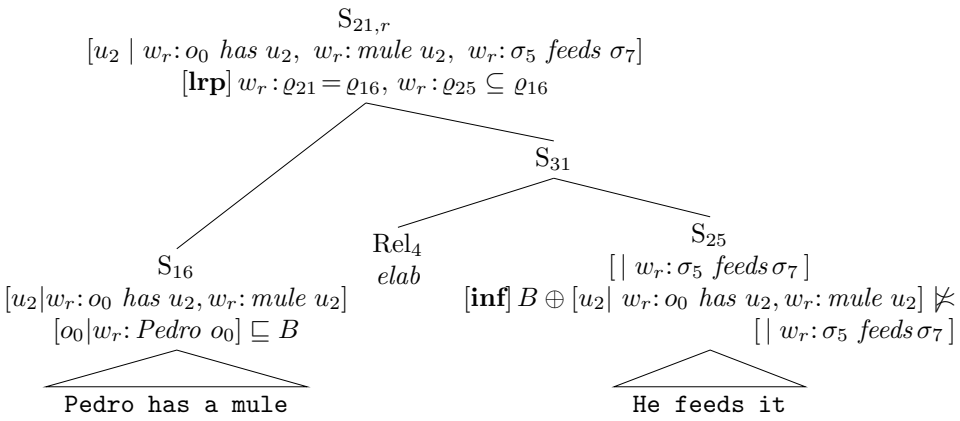

Ignoring the spatiotemporal and mereological meaning components of the elaboration, indicated by [**lrp**], for the sake of simplicity, the semantic value of the root of the described tree structure can be computed to be

$$\sigma_r = [\ u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7]$$

This represents what is asserted in the discourse. As was argued in the preceding chapter, discourse meanings consist of an implicit background B updated with σ_r. The discourse meaning that can be inferred from (5.4) is an underspecified DRS $B \oplus [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7]$, where the constraints triggered by the context-sensitive elements in the discourse must be satisfied. Two such constraints can be collected from the discourse description; one is triggered by the name 'Pedro' (5.5a), and the other is the informativity requirement triggered by the discourse relation (5.5b). We will not zoom in on the specific impact of these constraints here; look back to section 4.4.6 in chapter 4 for a discussion of their contribution to discourse meaning.

(5.5)a. $[o_0 \mid w_r{:}\, \mathit{Pedro}\ o_0] \sqsubseteq B$

b. $B \oplus [u_2 \mid w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2] \not\models [\,\mid w_r{:}\, \sigma_5\ \mathit{feeds}\, \sigma_7\,]$

Given (5.5a), axiom $\mathcal{A}19$, and idempotency of $\oplus$, the discourse meaning that can be inferred from (5.4) can be strengthened to

$$B \oplus [u_2\ o_0\ w_r \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7]$$

On the basis of the properness condition $\mathcal{A}21$ it can be inferred from this that

$$\mathit{proper}(B \oplus [u_2\ o_0\ w_r \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\, \sigma_7\,])$$

It follows from this that σ_5 and σ_7 are present in the universe of the discourse meaning. Given the definition of properness, the language user may infer

$$[\sigma_5 \mid\,], [\sigma_7 \mid\,] \sqsubseteq\ B \oplus [w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2, w_r{:}\, \mathit{mule}\ u_2]$$

and conclude for σ_5 that $\sigma_5 \in \{w_r, o_0, u_2\}$, or that σ_5 is in the universe of B but not in $\{w_r, o_0, u_2\}$. Similarly, for σ_7 he may infer that $\sigma_7 \in \{w_r, o_0, u_2\}$, or that σ_7 is in the universe of B but not in $\{w_r, o_0, u_2\}$. The same set of potential readings results as in our discussion of this example in the preceding chapter, section 4.4.7. In each case, the binder may be a discourse marker introduced earlier in the discourse, or it may be an entity in the universe of the implicit background of the discourse. In fact, as we will argue in section 5.4, pronouns in addition require binders that count as most salient given the ongoing discourse or occurrences in the nonlinguistic context. Hence, if the binder is an entity in the universe of the implicit background of the conversation, it must be a salient one. This is generally only the case when it represents an entity which is present

in the nonlinguistic context of interpretation and can be pointed out directly in the state of evaluation of the discourse.[5] This implies that the pronoun receives a deictic interpretation.

As before, various ungrammatical resolutions can be excluded by adding sortal constraints on discourse referents, and constraints on number and gender agreement. If this is done, four possible readings of the discourse remain. The following alternative specifications of information in $B \oplus \sigma_r$ indicate these readings.

(5.6)a. $[w_r\, o_0\, u_2 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, o_0\ \mathit{feeds}\ u_2]$
$\sqsubseteq B \oplus \sigma_r$

b. $[w_r\, o_0\, u_2\, \sigma_7 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, o_0\ \mathit{feeds}\ \sigma_7]$
$\sqsubseteq B \oplus \sigma_r \wedge [\sigma_7 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ]$

c. $[w_r\, o_0\, u_2\, \sigma_5 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ u_2]$
$\sqsubseteq B \oplus \sigma_r \wedge [\sigma_5 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ]$

d. $[w_r\, o_0\, u_2\, \sigma_5\, \sigma_7 \mid w_r{:}\, \mathit{Pedro}\ o_0,\ w_r{:}\, o_0\ \mathit{has}\ u_2,\ w_r{:}\, \mathit{mule}\ u_2,\ w_r{:}\, \sigma_5\ \mathit{feeds}\ \sigma_7]$
$\sqsubseteq B \oplus \sigma_r \wedge [\sigma_5 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ] \wedge [\sigma_7 \mid\ ] \not\sqsubseteq [w_r\, o_0\, u_2 \mid\ ]$

The first possibility corresponds to the preferred reading (Pedro feeds the donkey), the second to the case where *it* was taken deictically, the third to a deictic reading of *he*, and the fourth to the case where both pronouns pick up an extralinguistic referent. The hearer must now make a choice as to what was meant on the basis of his world knowledge and of what he knows or may assume is in the background.

5.2.2 Semantic Accessibility and Parallelism

Pronouns carry underspecified discourse markers as their semantic values. Through the interaction with the properness condition on discourse meanings and local contexts, it follows that these anaphoric discourse markers are bound in the pronoun's local context. Given the definitions of 'free' and 'proper' and the composition of discourse meanings and local contexts in our grammar, a binder must be accessible from the anaphoric marker (or from the pronoun) in the sense of standard DRT. We call this *semantic accessibility*. We will have, for example, that a discourse marker declared in the scope of a negation, implication, or disjunction cannot bind a discourse marker outside the scope of the operator: it is not semantically accessible. This accounts for well-known judgements such as the following.

(5.7) a. *John owns a*^{1} *car. It*$_1$*'s in the garage.*

[5] Remember we distinguish between o-markers and u-markers. o-markers are interpreted directly relative to i_0, the state of affairs the discourse is taken to be about, resulting in referential or deictic interpretations of nominal anaphors, depending on the type of anaphor. Due to $\mathcal{A}19'$: $\neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$ the universe of the implicit background contains only o-markers. See section 6.2 for some further discussion of this, and a proposal for an alternative treatment which relaxes the constraints on markers in B somewhat.

b. # *John does not have a*[1] *car. It*$_1$*'s in the garage.*

c. # *Every*[1] *student discusses a*[2] *paper. He*$_1$ *misunderstands it*$_2$*.*

All anaphoric elements which introduce free discourse markers as (part of) their semantic value are predicted to be bound by some semantically accessible binder in their local context, or induce the accommodation of such a binder, or else fail to resolve. Interestingly, there are cases of pronoun resolution which seem to refute this and resolve to a semantically inaccessible binder. This typically occurs with pronouns in clauses or sentences that are semantically parallel to some contextually available piece of information. The following discourse illustrates the point; the relevant pronoun is underlined.[6]

(5.9) a. *If Molly sees a stray cat, she pets it.*

b. *But if Dan sees <u>it</u>, he takes it home.*

Asher and Lascarides (2003), ex.(25), p.152

The connective 'but' in (5.9b) requires a relation of contrast, and hence, of parallelism, linking the two conditional statements. In line with Asher and Lascarides (2003) we assume that the coordinating connective 'but' attaches at top level, that is, it takes the whole of the preceding conditional as its left argument.[7]

In its most preferred interpretation—the so-called *sloppy* reading of (5.9b)—the right argument of 'but' means "if Dan sees a stray cat, he takes it home." Crucially, this is not a coreferential reading: the cat Dan sees need not be the one Molly sees, whenever she sees one. Rather than picking up the discourse referent introduced through 'a stray cat' in (5.9a), the pronoun 'it' in 'If Dan sees it' resolves to the parallel nominal description 'a stray cat'. It could be claimed that this violates semantic accessibility, since the parallel material that figures as the antecedent is in the scope of the conditional operator. We would suggest, however, that the pronoun in (5.9b) does not resolve to this antecedent through 'serial' binding; it resolves through parallelism and results in a form of analogy or counterpartness, rather than coreference. As long as the resulting sentence

[6]The phenomenon is well-known and also goes under the name of *pay-check pronouns*, after Karttunen's famous example:

(5.8) *The man who gave his paycheck to his wife was wiser than the man who gave it to his mistress.* Karttunen (1969)

[7]Both SDRT and our discourse grammar also allow for attachment at a lower position, so that the connective takes the consequent of the conditional as its left argument and the whole of the second conditional ends up in the scope of the first. This attachment would change the local context of the pronoun, and make available the discourse marker introduced by 'a stray cat' as a binder. This allows for a *strict* reading of the pronoun, but would also result in an overall different semantic value for the discourse than with an attachment at top level. The fact that this is possible does not obviate the need for some additional mechanism such as the computation of parallelism, to account for sloppy readings of pronouns.

or discourse meaning respects binding requirements there is no problem with semantic accessibility.

Exactly how pronoun resolution through parallelism should be treated in LDG is a question we will leave open. Let it suffice to observe that under certain conditions, personal pronouns may receive a full descriptive interpretation, rather than a bound variable one.[8] Furthermore, it is clear that the computation of parallellism plays a central role in the resolution process. In fact, parallellism per se (Prüst 1992; Gardent 1997b), or as constrained by focus interpretation (Gardent and Kohlhase 1995; Gardent and Kohlhase 1996; Pulman 1997), or discourse relations (Asher 1995; Asher and Lascarides 2003), affects the resolution of all types of anaphoric constructs, including those that are not referential in nature, such as VP-ellipsis. One of the desiderata of our discourse theory is to fill in this parameter.

5.2.3 Kataphors

The properness requirement $\mathcal{A}21$ enforces that every anaphor has a binder in its local context. This leads to a problem with the analysis of so-called anticipatory anaphors or *kataphors*, if discourse is processed strictly on a clause by clause basis. Kataphors are anaphoric elements which find their antecedent in the subsequent, rather than the preceding discourse. The following example illustrates this. The kataphoric pronoun 'them' is followed by its antecedent 'the nicotine plasters'.

(5.10) *Jack*[1] *has been trying to stop smoking for ages.*
He$_1$ *never expected them*$_2$ *to work,*
but the[2] *nicotine plasters actually helped quite a lot.*

Since local context are defined in terms of the preceding discourse, just after the second clause in (5.10) has been processed the local context of the pronoun 'them' does not contain a suitable binder. The intended binder, introduced by 'the nicotine plasters' is only made availabe in subsequent discourse. Since we fail to come up with a suitable binder, the discourse is predicted to be ungrammatical at this point.

The simplest solution to this problem is to allow a form of 'lazy interpretation':[9] a pronoun may remain unresolved untill the next processing step has

[8]Building on Cooper (1979), Gardent et al. (1996) argue that pronouns are ambiguous between a 'bound variable' semantics and a 'full description' semantics, and that in the context of semantic parallelism, the second type of meaning can be selected. The computation of parallelism is formally modeled by means of higher order unification with equivalences.

[9]A suggestion which is also made in a brief discussion of kataphors in Kamp and Reyle (1993), pp. 163–165. Furthermore, an extensive implementation of the idea is put forward in Van Deemter (1990) who proposes 'patient DRT' for the treatment of kataphors. The main idea is that an anticipatory anaphor introduces an *incomplete* discourse entity, to be completed by subsequent material under certain conditions. This converges with our own suggestion—further

taken place. This clearly helps us out in the case of (5.10), where the processing and interpretation of the third clause leads to the introduction of a suitable binder in the local context of the pronoun. As we shall argue in section 6.2, this is the result of accommodation, induced by the definite description 'the nicotine plasters'. Nothing needs to be changed about the properness requirement. Interestingly, this account predicts that the antecedents of kataphors are presuppositional expressions, names, and specific indefinites, or in general elements with a certain accommodation potential, since otherwise the binders will not be semantically accessible to the pronouns in their local context.

A quick at look at the overview of anticipatory anaphora in English in Huddleston and Pullum (2002), however, shows that while this is certainly a strong trend, it not always true, at least not for non-nominal antecedents:

(5.11) a. *It$_1$'s ridiculous! They've given the job to Pat*[1].

b. *Even though it$_1$'s not absolutely certain, there are very good grounds for believing that Ann was responsible for the change of policy*[1].

So a modification or refinement of our treatment shall be necessary. One option would be to relax the properness requirement on local contexts and, for instance, let it apply only to local contexts of completed sentences, rather than to contexts of all clauses. But this still doesn't account for cases like (5.11a).

Another option would be to allow the 'provisional' accommodation of an underspecified binder of a pronoun under condition that it will be fully determined as a result of the next processing step, and then counts as a salient topical element. After the next processing step, the provisionally accommodated binder can be identified with the discourse marker introduced in the domain of topical elements through the antecedent. This would explain why kataphors so often occur in clauses or discourse units which, in terms of the topic hierarchy of discourse, are subordinated to the discourse units that contains the antecedent. There would be no need to relax the properness requirement. The solution would have the added advantage of avoiding stages in processing the discourse where the discourse description is uninterpretable, at least where anaphoric reference is concerned.

We will not make a choice here, and leave as a matter for further research exactly how the constraints in our grammar should be set to fit in anticipatory anaphora. Relatively much is known about syntactic constraints on kataphora at sentence level. Further investigation would have to show how this interacts with discourse parameters such as topicality.

5.2.4 Summing up

We introduced the Properness requirement on discourse meaning and local contexts. Furthermore, we assumed properness to be a built-in property of our object

on—that the 'provisional' accommodation of a binder is allowed under certain conditions.

level relation of entailment. As a result, anaphoric elements which introduce free discourse markers as (part of) their semantic value are predicted to be bound by some semantically accessible binder in their local context. We touched upon two special cases of anaphoricity, namely so-called sloppy readings of pronouns, and kataphors. In the first case, it was argued that the pronoun 'resolves' through parallelism with its antecedent, but is not bound by it. Hence there is no violation of the Properness requirement. In the second case, it was suggested we adopt a form of lazy interpretation. As it turns out, the Properness requirement then still results in somewhat too strict predictions. Refinements of the discourse theory that would remedy this were touched upon.

As before, pronouns introduce underspecified discours markers, and the language user may infer potential binding relations from the discourse description and his grammatical knowledge. These potential binding relations correspond to the *possible solutions* of the resolution process in the sense of Van der Sandt (1992). They define, in our terminology, a set of *grammatical readings* of the discourse. All constraints in the grammar, either syntactic or semantic, potentially cut down the number of grammatical readings. In practice the language hearer will not consider these as equally acceptable readings. On the basis of his world knowledge and general preferences he will select a single or a few most plausible ones.

In the following sections we zoom in on two classes of further constraints in the grammar. In section 5.3, we put forward an implementation of Chomsky's Binding Principles. The Binding Principles characterise the individual binding behaviour of particular classes of anaphors, and can be incorporated in the grammar through syntactic constraints in the lexicon. In section 5.4 we briefly turn to constraints on pronoun resolution induced by salience and the topic structure of discourse.

5.3 Syntactic Binding Principles in LDG

Not every semantically accessible discourse marker is a possible binder for each type of anaphor. As is well-known, anaphoric expressions may differ with respect to the syntactic domain within which their binder can be generated. For instance, personal pronouns and reflexives exhibit very different binding behaviour. Although there are certain exceptions to the rule, it is generally the case that pronouns and reflexives occur in complementary distribution: if a given binder binds a pronoun, it cannot bind a reflexive in the same position, and vice versa. The following examples illustrate this.

(5.12) a. *Every*^{1} *handsome man admires himself*$_1$.
b. # *Every*^{1} *handsome man admires him*$_1$.
c. # *John*^{1} *believes that Mary hurt himself*$_1$.
d. *John*^{1} *believes that Mary hurt him*$_1$.

Intuitively, what seems to be the case is that when the binder is 'close by', for example, when it is an argument of the same predicate as the anaphor, the anaphor must be a reflexive, and cannot be a pronoun. This intuition is what underlies the Binding Conditions of Government and Binding Theory (Reinhart 1976; Chomsky 1981), which, in one formulation or another, require that pronouns are ^{GB}*free* within the relevant syntactic domain (their 'governing category'), while reflexives must be ^{GB}*bound* within the same domain. We tagged these terms ^{GB}*(term)* in order to distinguish them from the notions 'free' and 'bound' employed in the semantic theory of our grammar elsewhere. They refer to the following tree geometric notions.

(1) a. Node k *c-commands* node k' if every branching node dominating k dominates k'.

b. A node is ^{GB}*bound* if it is coindexed with a c-commanding node.

c. A node is ^{GB}*free* if it is not bound.

It is clear that these notions can be made available in our grammar if we read 'k and k' are coindexed' as $\sigma^{\pi}(k) = \sigma^{\pi}(k')$. The Binding Theory itself consists of the following three principles.[10]

(A) A reflexive or reciprocal must be ^{GB}bound in its governing category.

(B) A (non-reflexive) pronoun must be ^{GB}free in its governing category.

(C) An R-expression must be ^{GB}free everywhere.

R-expressions are DPs that are not pronouns, e.g. names, definite descriptions, and quantifiers. The syntactic literature contains much discussion about the correct definition of 'governing category', but a rough approximation is that the governing category of a node k is the lowest k' properly dominating k that is labelled DP or S. Using $\mathsf{gvc}(k)$ to represent the governing category of a node k, and $\mathsf{cc}(k, k')$ to denote that a node k c-commands a node k', principles (A), (B), and (C) can be formalized as additional constraints within the present theory, as follows.

[10] To avoid confusion with respect to the term 'anaphor', we use the formulation of the binding principles as given in Haegeman (1994). The original formulation in Chomsky (1981) is

(A) An anaphor must be ^{GB}bound in its governing category.

(B) A pronominal must be ^{GB}free in its governing category.

(C) An R-expression must be ^{GB}free.

Here 'anaphor' should be read as 'reflexive or reciprocal', while the category of 'pronominals' includes non-reflexive pronouns and 'R-expressions' are DPs that are not pronouns.

Binding Conditions in LDG

(A) $\forall k[\, \mathit{refl}(k) \rightarrow \exists k'(\mathsf{gvc}(k) \lhd^* k' \wedge \mathsf{cc}(k', k) \ \wedge \ \sigma(k) = \sigma(k'))\,]$

(B) $\forall k[\, \mathit{pron}(k) \rightarrow \neg\exists k'(\mathsf{gvc}(k) \lhd^* k' \wedge \mathsf{cc}(k', k) \wedge \sigma(k) = \sigma(k'))\,]$

(C) $\forall k[\, \mathit{pn}(k) \vee \mathit{det}(k) \rightarrow \neg\exists k'(\mathsf{cc}(k', k) \wedge \sigma(k) = \sigma(k'))\,]$

This formulation builds on a treatment of quantifiers in which either quantifiers are raised in syntax or logical form leaving a trace of the right semantic type, or the quantifier is 'raised' in the semantics, as in the in situ treatment of quantifiers proposed in Muskens (2001). In the latter case, the semantic value of the DP is of the right type, namely type π. The binding conditions constrain the semantic values of the traces left by quantifier raising, or, in the case of the in situ treatment of quantifiers that we adopt here, the semantic values of the DP's involved.

As the reader may check for himself the binding conditions (A) and (B) explain the grammaticality judgements in (5.12). Furthermore, condition (C) accounts for the ungrammaticality of cases such as (5.13a), while allowing coindexations as in (5.13b).

(5.13) a. # *He*$_1$ *admires Joe*[1].

b. *John*[2] *is very proud of his*$^1_2$ *brother, but I don't like the*$_1$ *boy.*

The Binding Conditions describe properties of certain word classes, and as such could also be integrated in the elementary descriptions of elements of these word classes. When a lexeme selects an elementary description of a given word class, the relevant binding condition is part of its lexical description. For instance, we would have the following lexical descriptions for the lexemes 'he', 'himself', 'Joe', and 'the', selecting the word classes *pron*, *refl*, *pn*, and *det* respectively.

(5.14)

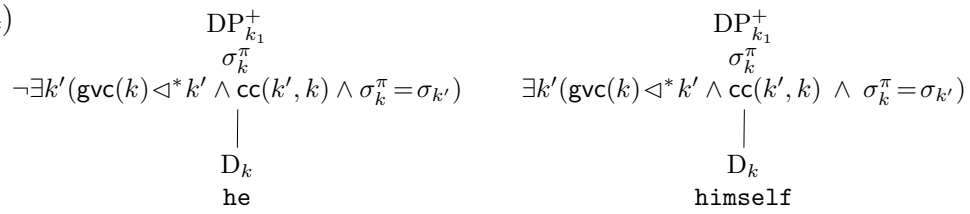

(5.15)

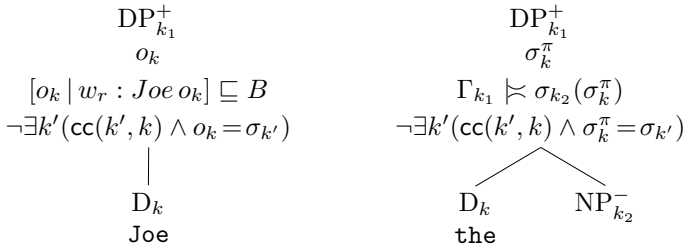

The added constraints characterise the syntactic domain in which the anaphor can, or cannot find a binder. These constraints of course interact with the general

requirements on properness formulated earlier in filtering out ungrammatical resolutions, and with the other characterising conditions attached to the anaphor. The lexical descriptions of 'Joe' and 'the' illustrate the latter. Beside the syntactic binding condition, each carries a semantic condition constraining its relation to the context of interpretation. In the case of proper names that condition ensures that the name denotes an element in the implicit background of the discourse, see section 4.4.7 of the preceding chapter. In the case of the definite determiner it expresses the presuppositional character of definite descriptions. We will zoom in on this in the second half of this chapter.

Note that the constraints on anaphora resolution discussed so far are all stated in a single describing logic, even though they came from syntax and semantics, two separate levels of the grammar. The description approach allows us to integrate theories from various levels and a syntactic theory such as the Binding Theory, if adopted, has immediate relevance for the construction of discourse meanings. In the next section we turn to the constraining effect of discourse topic structure and salience on pronoun resolution. While what we have to say there is largely programmatic, it does suggest that what, for want of a better term, may be called the *information status* of contextual information, can be captured in the grammar in terms of the syntactic and semantic entities and properties at our disposal. Thus, we expect that this parameter of the discourse theory can also be fully integrated in the grammar.

5.4 Discourse Topic Structure and Salience

Anaphors must be bound in their context of interpretation, but they can only be bound by a semantically accessible binder. As it stands, our grammar does not contain any structural requirement on the relation between the antecedent introducing the binder and the anaphor, apart from the Binding Conditions in the preceding section. It is, however, a common assumption in frameworks of discourse interpretation that, directly or indirectly, anaphora resolution is sensitive to discourse structure, and that various classes of anaphors must find their antecedent, or the binder it introduces, 'at the right frontier' of the discourse parse tree. To put it differently, the binder must not only be semantically accessible but also, in a sense yet to be made precise, *discourse available*.

In this section we want to make the notion of discourse availability a bit clearer. The matter is complicated for various reasons. First and foremost, discourse availability is not a simple structural notion: there is a handful of possibly overlapping factors, such as discourse topic structure, salience, focus of attention, back- and foreground structure, and the syntactic position of antecedents, the amalgamate of which is somehow responsible for making some binders available at a given structural position, and others not. We shall not investigate each of these factors. Our concern is the overall organisation of the grammar with dis-

course availability playing a subordinate role. The notion must be captured in terms of a combination of semantic factors, and these, like other parameters in our grammar, can be grafted onto discourse syntactic structure.

Different classes of anaphors come with different demands on the discourse availability of their binders. So, for instance, observations made relative to personal pronouns do not necessarily carry over to definite descriptions; whereas the resolution of VP-ellipsis is sensitive to discourse structure in a relatively straightforward manner, this is less obvious for personal pronouns; abstract object anaphors and demonstratives come with their own requirements yet again. Thus, we must be careful to distinguish the characteristics of specific classes of anaphors from general aspects of discourse availability.

Finally, the accounts or treatments offered in different frameworks are hard to compare whenever they presuppose different discourse structures. Explanations made in frameworks that presuppose the flexible structures of RST, notably SDRT, do not necessarily, at least not directly, carry over to frameworks which presuppose strict tree structures, such as LDM, LTAG for discourse, and LDG (in its current formulation). What is 'the right frontier' in the latter approaches is simply not the same thing as what it is in SDRT. In what follows, strict tree structures will be presupposed; when we discuss other structures this will be mentioned.

Let us zoom in on personal pronouns, and consider in what sense their binders must be discourse available. Consider the following example, taken from Asher and Lascarides (2003) pp. 8–10. The example was already briefly discussed in chapter 3, section 3.4.

(5.16) a. *Max had a great evening last night.*
b. *He had a great meal.*
c. *He ate salmon.*
d. *It was a beautiful pink.*

The pronoun 'it' in sentence (5.16d) felicitously refers to the salmon Max ate, introduced in the context of interpretation through (5.16c), the immediately preceding utterance. The statement 'it was a beautiful pink' elaborates on 'he ate salmon', zooming in on a property of the salmon. Now compare

(5.17) a. *Max had a great evening last night.*
b. *He had a great meal.*
c. *He ate salmon.*
d. *He devoured lots of cheese.*
e. *He then won a dancing competition.*
f. *# It was a beautiful pink.*

Given background knowledge about what sort of things are normally pink, the

most likely interpretation of (5.17f) is one in which the utterance elaborates on (5.17c) again: the salmon Max ate was a beatiful pink. The anaphoric link between the pronoun in (5.17f) and the antecedent in (5.17c) is infelicitous, however. The reason for this, discourse theory claims, is that the binder which is introduced in the domain of discourse by the antecedent, though it is semantically accessible to the pronoun, is not discourse available to it.

So when is a binder discourse available to a given anaphor? As a first try, we might say that the S-node or the minimal discourse unit at which the binder can be assumed to be introduced or generated must be a potential attachment site of the tree structure described by the input sentence description that contains the anaphor. It must be a node 'at the right frontier' of the tree structure described by the input discourse description. In the case of (5.17f) it isn't. To see his, consider the following picture, which is the discourse description of (5.17), as it can be taken to result just after the incrementation step induced by (5.17f).

(5.18) Δ_6:

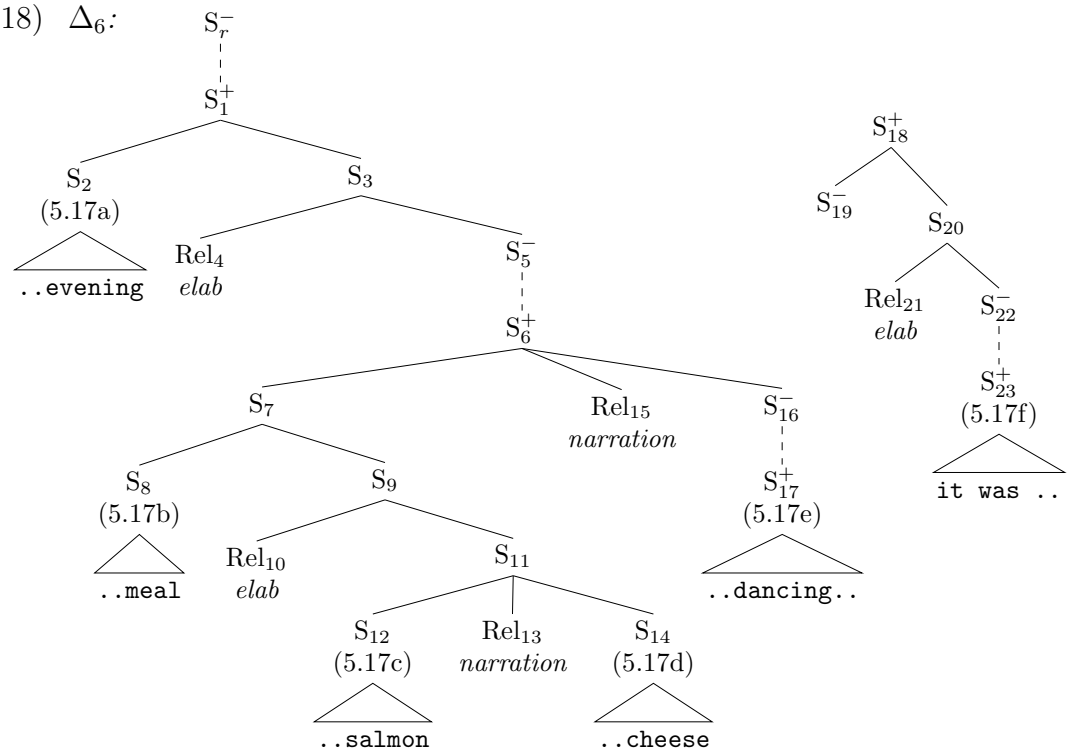

The discourse context of (5.17f) is represented by the part of the discourse description in (5.18) that describes (5.17a) to (5.17e). The right frontier of this part of the discourse description is the set of nodes that may serve as an attachment site for subsequent updates: the nodes named r, 1, 5, 6, 16, and 17. The discourse unit at which the binder of the pronoun in (5.17f) may be assumed to be introduced is node 12, but this is not a potential attachment site for the tree structure described by the input sentence description, containing the pronoun. Hence, the binder is not discourse available to the pronoun.

Setting aside for the time being the obvious question what it means for a discourse unit 'to generate or introduce a binder', this looks like an acceptable explanation. As the following discourse illustrates, however, we run into problems with the assumption that the discourse unit introducing a potential binder must be one of the structurally accessible nodes of the discourse context. Consider

(5.19) a. A^1 *man was run over.*
b. *The bus was going too fast.*
c. *The driver didn't see him*$_1$.

Asher and Lascarides (2003), ex.(24), p.151

The utterances (5.19b) and (5.19c) together function as an explanation of the fact conveyed through (5.19a). The discourse description resulting after processing the first two utterances may be graphically represented as follows.

(5.20) Δ_2:

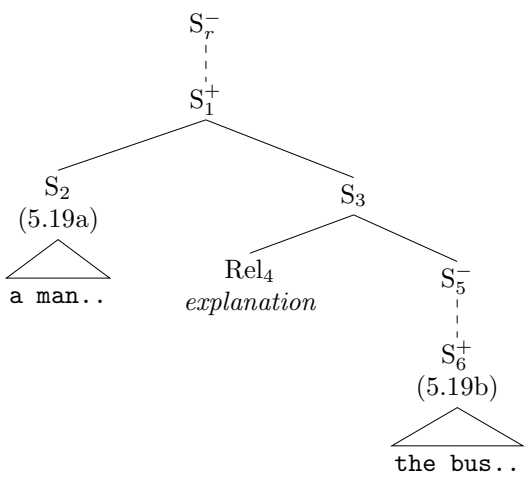

The discourse description resulting from the next processing step, and some subsequent reasoning about discourse relations, may be given as follows.

(5.21) Δ_3:

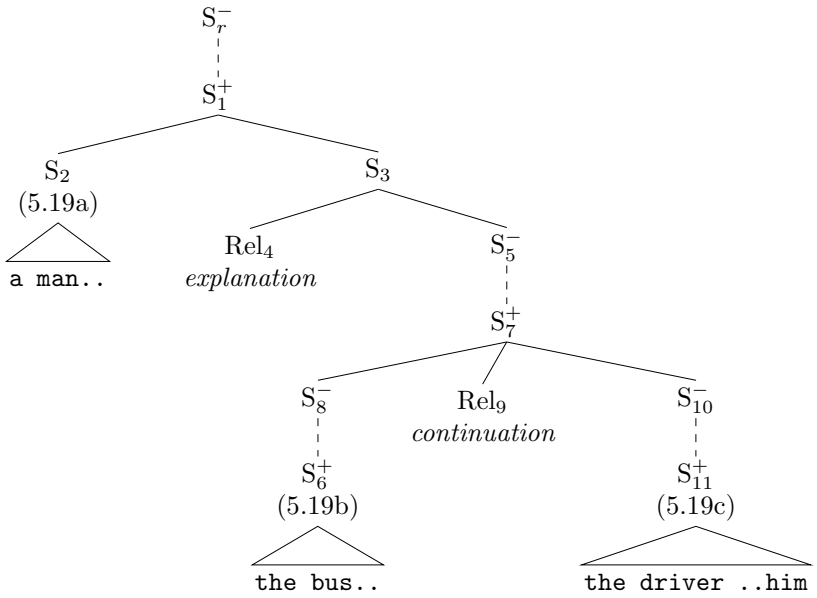

The pronoun in (5.19c) resolves felicitously to a binder which is introduced at a discourse unit higher up in the embedding discourse context, namely the one represented by node 2. However, this is not one of the possible attachment sites of the update that was induced by (5.19c), namely the nodes named r, 1, 5, and 6. It is incorrectly predicted that the binder introduced by 'a man' is not discourse available. We conclude that, given the type of discourse syntactic structures described by our grammar, discourse availability cannot be defined directly, or uniquely in terms of structural accessibility, as it is defined in our grammar system.

This does not imply that structural relations in general are not relevant. The reader may have noticed that there is a relation of c-command between node 2 and node 11 in verifying trees of Δ_3. So perhaps the relation between antecedent and pronoun can be captured by requiring that a pronoun must be *c-commanded* by the discourse unit that introduces its binder. The overall treatment of anaphora resolution in for instance the Linguistic Discourse Model and in SDRT suggests that c-command on its own will not be enough, however. There, the difference between coordinating and subordinating discourse relations is crucial. Transporting this insight to our own grammar system, this implies that the binder introduced at node 2 is available not only because node 2 c-commands node 11, but also because it 'subordinates' it in a certain sense.

Let us zoom in on the treatment provided in SDRT. SDRT presupposes a more flexible type of discourse structures than simple tree structures. The hierarchical structure of SDRSs emerges in part on the basis of outscoping relations on discourse constituents, comparable to the subordination relation on DRSs in standard DRT. However, it also depends on the relation of *topic subordination* implied by certain discourse/rhetorical relations. For example, the explanation and the elaboration relation induce a topic subordination relation between their arguments: the left argument is a topic for the right argument. In the resulting

hierarchical structure the effect of this that, in our own words, the left argument 'dominates' the subordinated right argument. In the example under consideration, this implies that the discourse segment representing the left argument of the explanation (5.19a) is at the right frontier of the topic hierarchical structure. Hence, the binder 'a man' is predicted to be discourse available to the pronoun.

In principle, the LDG formalism could be modified so that the 'right' nodes end up at the right frontier of the discourse structure. This is possible if we relax the type of graph structures that can be described by discourse descriptions in LDG. A treatment similar to SDRT may then be obtained. We prefer to have simple tree structures, however. The essential insight is that discourse availability is determined, at least in part, by the discourse topic structure. A discourse topic structure may be defined on the basis of the specific characteristics of discourse relations and connectives, the information structure of the input clauses,[11] and the syntactic composition of the discourse per se. An independent mechanism could be defined, that walks through the embedding discourse context top down and from left to right and delivers at each node of the discourse tree a set of binders that count as *locally topical* given the topic hierarchy of the discourse context. Together with other factors this defines what counts as a discourse available binder at a given node.

What would be those factors? We tentatively group together the salience of the binder, its level of activation, and the measure to which it is in the focus of attention. All of these center around a cognitive, gradient notion of salience or topicality, applied to distinct entities in the domain of discourse. See for instance Ariel (1990) for a useful discussion of anaphora, salience and accessibility levels from a cognitive linguistic perspective, and Grosz et al. (1995) and its precursor Sidner (1986) for a treatment of focus of attention in Centering Theory.[12] Broadly speaking, the topic hierarchy of discourse and salience effects determine the discourse availability of potential binders. It might even be that these two parameters can be fused into one notion of 'information status' but we leave this as a matter for further investigation.

For concreteness, we will assume that salience results in an ordering over the discourse markers that count as topical at a given node in the discourse structure given the topic hierarchy of the discourse. Different classes of anaphors may then

[11] We use the term *information structure* to refer to the division of clause meanings into parts with a distinct update function relative to the context of interpretation. In much of the relevant literature, two basic types of structures are distinguished: first, a division into background (or focus frame), one or more foci, and contrastive topics, and second, a division into topic and comment. Givenness and contrast constraints can be projected from the first type of structure, constraints on subordinating topic relations and common topics can be projected from the second. For an overview of topic-comment research, see Van Kuppevelt (1997). For an entrance to topic-focus literature and focus interpretation theory, see chapter 3 of Vallduví (1990), and Von Stechow (1991).

[12] For a computational treatment of focus of attention based on Sidner (1986), see Blackburn and Bos (1999).

be distinguished according to their sensitivity to topicality per se, and to the level of salience of topical elements. In line with Grosz et al. (1983), Grosz et al. (1995), pronouns may be assumed to select topical binders that are topmost in the salience order. Using some completely programmatic but suggestive notation, we may indicate how a requirement of this kind can in principle be incorporated in the elementary description of pronouns:

(5.22) *pron:*

$$\begin{array}{c} \mathrm{DP}^{+}_{k_1} \\ \sigma_k \\ \sigma_k \in T^{mostsalient}_{k_1} \\ | \\ \mathrm{D}^{\Diamond}_{k} \end{array}$$

In the case of example (5.19), with discourse description (5.21), the binder introduced by 'a man' at node 2 should be in the set of most salient topical elements at the node representing the maximal projection of the pronoun 'him', dominated by node 11. This would be because the *explanation* relation sets its left argument (the fact or eventuality to be explained) as a main topic for the right argument (the explanation), and hence carried over as a common topic of the *continuation* to node 11. This would make, among others, the binder introduced by 'a man' into a topical element at node 11. For being the patient or theme of the running over eventuality it might be taken to count as a most salient topical element at node 11. As such it may be inherited by the maximal projection of the pronoun and count as a discourse available binder. In contrast, in the case of our earlier example (5.17) with discourse description (5.18), the potential binder introduced by 'salmon' would count as a topical element at node 12 given the discourse topic hierarchy, but not at any of the possible attachment sites for node 19. Hence it cannot be a topical element at the maximal projection of the pronoun dominated by node 23, whatever the attachment of the update will be.

For other classes of anaphoric elements constraints on discourse availability of the binder may be different. For example, it is often assumed that definite description do not require salient antecedents or binders. Given the distinctions introduced above we then have a choice between requiring the binders of definite descriptions to be topical but not salient, or not even requiring them to be topical. The two options are illustrated in (5.23).

(5.23) *definite det:*

$$\begin{array}{c} \mathrm{DP}^{+}_{k_1} \\ \sigma^{\pi}_{k} \\ \sigma_k \in T_{k_1} \\ \Gamma_{k_1} \models\!\!\succ \sigma_{k_2}(\sigma^{\pi}_{k}) \\ \mathrm{D}_k \qquad \mathrm{NP}^{-}_{k_2} \\ \texttt{the} \end{array} \qquad \begin{array}{c} \mathrm{DP}^{+}_{k_1} \\ \sigma^{\pi}_{k} \\ \Gamma_{k_1} \models\!\!\succ \sigma_{k_2}(\sigma^{\pi}_{k}) \\ \mathrm{D}_k \qquad \mathrm{NP}^{-}_{k_2} \\ \texttt{the} \end{array}$$

Given the left lexical description, the binder of a definite description must be in the domain of topical elements determined by the topic hierarchy, given the right

description, the choice of a binder is completely free, apart from other constraints. Options such as these may guide a further investigation of what notions such as 'being a topical element given the topic hierarchy of the discourse' and 'salience' stand for. For a glimpse at the search space, consider the following data.

(5.24) a. *Max had a great evening last night.*
b. *He had a great meal.*
c. *He ate salmon.*
d. *It / the salmon was a beautiful pink.*

(5.25) a. *Max had a great evening last night.*
b. *He had a great meal.*
c. *He ate salmon.*
d. *He devoured lots of cheese.*
e. *# It / # the salmon / the SALMON was a beautiful pink.*

(5.26) a. *Max had a great evening last night.*
b. *He had a great meal.*
c. *He ate salmon.*
d. *He devoured lots of cheese.*
e. *He then won a dancing competition.*
f. *# It / # the salmon / # the SALMON was a beautiful pink.*

As the examples show, a definite description 'the salmon' replacing 'it' in the example under discussion can be used in more positions than the pronoun, though not in all. Furthermore, in some positions the definite description must be accentuated, so that it gets a subsectional or contrastive interpretation. These judgements may be accounted for if we assume that the binder introduced by 'salmon' at node 12 is topical only at node 12, and definite descriptions do not require topical binders. The accentuation of the definite description in (5.25) and (5.26) is there for independent reasons.

Another possibility would be to assume that definite descriptions do require topical antecedents and that in the case of (5.25) this requirement is satisfied by resolving subsectionally to an element that is topical, namely the meal at node 7 or the summation of what is eaten at node 11. To account for the infelicity of the accentuated definite description in (5.26) it could perhaps be argued that the binder introduced by 'a salmon' is too deeply entrenched in the topical elements at the potential attachment sites 1 and 6 to be reached by means of subsectional anaphora.

Yet another option would be to assume that the binder introduced by 'salmon' is not only topical at node 12, but also at node 11 and 7. The infelicity of the

pronoun in (5.25) could then be accounted for on the assumption that the binder, though topical, is not most salient at the nodes 11 and 7. Definite descriptions, however, do not require salient binders, so it is a discourse available binder to the definite description given attachment at node 11 or 7. The accentuation of the definite description may serve to indicate the unsalient status of the binder. To account for the infelicity of the accentuated definite description in (5.26) it could be assumed that definite descriptions require topical binders, but that the putative binder is not topical at the potential attachment sites node 1, 6, and 17.

We believe each of these options to be worth considering. We do have a bias towards a treatment of definite descriptions that makes them sensitive to what is the local discourse topic or what count as locally topical elements.[13] It seems unlikely that some form of 'anchoring' to the local discourse topic would be completely irrelevant for the resolution of definite descriptions. Since the nominal material can help to identify a suitable binder, the kind of link that a definite description might establish towards elements in the topic domain can be more flexible or complex than the links established by plain (singular) personal pronouns. Perhaps binders of definite descriptions are not themselves required to be topical, but must link up through bridging or inference to elements that are topical given the discourse topic hierarchy.

Concluding, various points need to be clarified before we can provide a full-fledged treatment of discourse availability. What sort of semantic objects are discourse topics? How do discourse relations, connectives and other operators in the lexicon constrain the discourse topic hierarchy, and how does the information structure of input clauses interact with this? How do we reason from discourse topics to local domains of topical elements? What does it mean for a discourse unit to introduce a binder? To what extent can binders be inferred or accommodated and thus end up in local topic domains without having been overtly introduced through an antecedent? What determines a salience ordering over topical elements? Can this be defined purely in terms of syntactic, semantic, and prosodic parameters of the grammar, or does it also involve incorporating cognitive linguistic or 'processing' constraints such as management of memory load in the grammar?[14] These questions are not new. Anaphora is a long-standing topic in linguistics, and it

[13] If only to make sense of the *unicity* or *exhaustification effect* that is commonly assumed to be associated with definite descriptions. The claim that definite descriptions refer uniquely to some element in the domain of discourse that satisfies its nominal description can only be maintained when relativised to a very restricted and dynamically changing domain of candidates. As pointed out in Lewis (1979), just the assumption that this is a contextually determined domain of discourse is not enough. In his words, "*the F* denotes x if and only if x is the most salient F in the domain of discourse, according to some contextually determined salience ranking." Since Lewis does not talk about topicality, we are not sure how his notion of 'most salient' elements compares to ours. Possibly it covers elements in the domain of discourse which we would simply call 'topical'.

[14] Compare for instance Ariel (1990), Walker (1996b).

might be called the cornerstone of discourse theory. Much is known about the empirical facts in this field. Organising and explaining the facts in terms of a fully transparant and unified theory of 'informational status' is still a challenge, however.

5.5 Accommodating Binders

Anaphora resolution is modelled as a reasoning process in LDG. Anaphoric elements occurring in a discourse introduce underspecified discourse markers,[15] and the language user's grammar tells him that these markers must have binders in the local context of the pronoun, as determined by the discourse context. Given the general constraints in the grammar, and the specific syntactic, semantic, and informational constraints that classify the anaphor, the hearer may reason about its possible resolutions. Given his knowledge of the actual utterance situation, his world knowledge, and his preference for, for instance, selecting most plausible background assumptions, he may determine a few ways to resolve a given anaphor, or a single most preferred one.

The reasoning process generally involves selecting possible models of the context of interpretation that fit the preferences of the hearer and the requirements of the grammar relative to the discourse under consideration. The context of interpretation, which is represented in our grammar by the implicit background B, is largely underspecified. The hearer may accept information in B in order to satisfy constraints triggered by the interpretation of the ongoing discourse. Thus, B gets specified, or as it were, is partly constructed, as a result of the interpretation process. This is *global* accommodation, corresponding to accommodation of the initial context of interpretation in possible worlds approaches, and to accommodation of the main DRS in DRT. A generalisation of this to local and intermediate accommodation shall be provided in the next chapter.

Here, we are concerned with the interaction between anaphora resolution and accommodation, in particular with what might be called the 'accommodation potential' of pronouns. It is a common observation that pronouns allow accommodation of their binders much less easily than, for instance, definite descriptions. Given what we have seen so far of pronouns in our grammar, this must be explained from the interaction of their individual syntactic and semantic properties with accommodation in general. The challenge for LDG is not so much to explain that accommodation occurs, but to make clear how overgeneration is avoided.

Obviously, the hearer's linguistic and nonlinguistic knowledge keep accommodation under control in the sense that whatever is accommodated, the constraints in his grammar and his preferences must be respected. There is a general agreement that assumptions to be accommodated are to be uncontroversial and un-

[15] Some anaphoric elements, for instance VP-ellipsis, gapping, and 'one', introduce underspecified *properties*. Parallelism requirements play a central role in resolving them.

surprising (Heim 1992, page 212). In fact, as argued in Stalnaker (1998), this can be explained in terms of an appropriateness requirement linked to the concept of speaker presupposition: the speaker presupposes that the addressee is willing to accept the presupposition. Our account stands in need, however, of a precise explanation of what *justifies* the acceptance of a given background assumption, or what counts as sufficient evidence for the hearer given the type of information that must be accommodated.

To illustrate the above, consider the following discourse. This is Partee's well-known 'ten marbles' example.[16]

(5.27) *I dropped ten marbles and found only nine of them.*
It's probably under the sofa.

The second sentence in (5.27) is unfelicitous, because the pronoun 'it' cannot very well be used to refer to the tenth marble. But if, as we say, pronoun resolution is a reasoning process, then why should a hearer not infer, on the basis of his world knowledge and the information provided in the first sentence of the discourse, that there is a marble, namely one of the ten dropped by the speaker but not found, accommodate that information in B, and assume that the pronoun is bound by it? This would result in a perfectly coherent discourse, in which the second sentence explains where the tenth marble might be. Why is accommodation of the binder blocked? Note that if we replace the pronoun by a definite description, the problem disappears:

(5.28) *I dropped ten marbles and found only nine of them.*
The tenth marble is probably under the sofa.

This suggest that inferrable information can in principle be accommodated, if the right kind of anaphor is used. We propose to explain the infelicitousness of the second sentence in (5.27) in terms of additional requirements on potential binders associated with pronouns. As was argued in section 5.4, pronouns, unlike definite descriptions, come with constraints on the informational status of their binder. The binder must be a salient topical element in the local context of the pronoun. It may achieve this status through linguistic properties of the preceding discourse, or through occurrences in the nonlinguistic context of the conversation. Just the fact that a binder can be inferred does not make it into a salient topical discourse entity, however. Crucially, the added requirement cannot be satisfied by simply *assuming* the inferrable binder to be salient and topical. Neither the nonlinguistic context of utterance nor the discourse context sufficiently justify this. Hence accommodation fails, or counts as strongly dispreferred.

[16]Compare Stalnaker (1998), where example (5.27) starts off a discussion about the way in which an abstract context might make individuals available for reference.

An account like this does not predict that accommodation of a binder is completely impossible with pronouns. Rather, it predicts that it only occurs when the relatively 'heavy' conditions on binders triggered by pronouns can be satisfied, or assumed to be so, given the hearer's linguistic and nonlinguistic knowledge and the discourse context established so far. To illustrate this, consider the following discourse.

(5.29) ? *We don't have a car. We sold it.*

The discourse marker introduced by 'a car' in the first sentence is in the scope of a negation and is therefore not semantically accessible from the pronoun in the second sentence. Still, in our view, there is a felicitous interpretation of the discourse in (5.29), which can be paraphrased as 'We don't have a car. We sold the car that we used to own.' According to our proposal, this reading results from accommodating that there is a car which the speaker used to own, and accommodating that that car is a most salient topical entity given the immediately preceding discourse.[17] This is possible, we suggest, if it can be taken to be common knowledge between the speaker and hearer of the discourse that the speaker owned a car, and the having of a car can be taken to be the discourse topic under discussion.[18] The first assumption is easily justified in a world where almost everybody owns a car, the second involves selecting an initial context of interpretation which is slightly primed.

The selection of a somewhat loaded initial context of interpretation occurs more often. Consider

(5.32) *He opened the door. The room was pitch-dark.*

The discourse describes a situation, starting in medias res, so to speak. Interpreting the first sentence, we accommodate without effort that there is some male individual who opened a door. We accept that this individual is a salient topical element in the narration that is going on, without any support from the

[17]The option that the car is a salient entity in the nonlinguistic context of utterance can be discarded as being extremely unlikely.

[18]Notice that this reading can be reinforced with the use of accentuation, in various ways. One option is to use accentuation that signals polarity focus/ 'verum-focus'/ focusing of the truth value in the first sentence, and VP focus in the second, as in (5.30).

(5.30) *We DON't have a car. We SOLD it.*

Another option is to use an accentuation that signals contrastive focus (expressing a contrast between having, and selling a car) as in (5.31).

(5.31) *We don't HAVE a car. We SOLD it.*

In each case, the articulation of the information structure of the sentences helps to accommodate a discourse topic, relative to which the use of the pronoun becomes completely felicitous.

immediately preceding discourse context. It has often been pointed out that this typically occurs with the opening sentence of a story, or in general, narrative discourse. Apparently, in that context, the opening sentence or sequence of sentences is allowed a kind of bootstrapping function. It may trigger the initialisation of the implicit background with, among more, the main characters of the story, on the basis of no other evidence than the use of the pronouns themselves. Apparently, the authority of the speaker, who supposedly knows what he is talking about, is sufficient justification for the accommodation of binders.

Summing up, as an underspecification formalism, LDG allows a natural treatment of global accommodation. Global accommodation is the partial specification of the largely underspecified general background B of the conversation, resulting from the need to satisfy constraints that are triggered by the interpretation of the ongoing discourse. As will be discussed in the next chapter this notion can be generalised to cover the specification of *local* background information, so that accommodation may also contribute to the construction of discourse semantic content. Accommodation necessarily respects the general constraints of the language user's grammar and his preferences. For instance, since discourse meanings must be consistent no information can be accepted in B that results in inconsistency of the discourse meaning. And since the language user prefers more plausible readings of a discourse meaning over less plausible ones, he will not accept information in B that results in an implausible discourse meaning, when more plausible alternatives are available.

Where anaphora resolution is concerned, the accommodation of binders is further restricted by the specific syntactic, semantic and informational conditions that characterise the anaphor's relation to its local context. Pronouns come with the requirement that their binder, apart from being semantically accessible, must be a salient, topical entity. It was proposed that the fact that the presence of a binder can be *inferred* from the local context is not sufficient to satisfy the salience requirement. Salience requires a different kind of justification or evidence from the discourse context. A salient binder is accommodated only if justification is available or can itself be accommodated without much cost. In future work, it needs to be spelled out what sort of information requires what sort of justification, and how the cost of accepting a piece of information depends on the available evidence. The theory of accommodation must be backed up with a treatment of justification.

5.6 Conclusion

We have not at all attempted to cover a vast range of anaphoric elements in LDG in this chapter. Neither did we propose an analysis of basic cases of anaphora that differs radically from what is commonly assumed. With a view to the application

of our discourse model to presupposition theory, to be unfolded in the next chapter, we were interested in capturing the discourse parameter of anaphoricity as such. We proposed to do so by letting anaphors be elements in the language that introduce underspecified discourse markers, and stating as a general axiom in the grammar that discourse meanings and local contexts must be proper. Given the composition of local contexts it can be shown to follow from this that anaphors must be bound in their local context.

The binding relation, though underspecified, is subject to all interacting constraints in the grammar. A benifit of using an underspecification mechanism is that all of this can be handled declaratively in the grammar. LDG shares this feature with UDRT and (U)SDRT. There is also a procedural side to our account, and this is in the manner in which *resolution* is modeled. For any given anaphor, the hearer must find out what its potential binders are, and, if possible, pinpoint a specific most preferred one. This is the outcome of a reasoning process in which both linguistic and nonlinguistic knowledge are involved. The architecture of our discourse model implements the idea that the hearer may access these knowledge sources in tandem.

The reasoning process may result in global accommodation: the partial specification of information in the implicit background of the discourse. In our approach, binders of anaphors can in principle be inferred and accommodated. This, as we shall see in the next chapter, is particularly attractive with a view to bridging anaphors. Characterising conditions associated with each type of anaphor constrain the accommodation process. In the case of pronouns it was suggested that their binders must be salient topical elements and that this is what makes accommodation of a binder for a pronoun costly—though not completely impossible. As will be argued in the next chapter, definite descriptions come with different requirements on their binders and therefore have a different 'accommodation potential'.

The treatment raises some interesting questions for further research. For one, if accommodation of background information can be more, or less costly, an interpreting agent must somehow be able to weigh these costs according to his grammatical knowledge and preferences. How does he decide whether the accommodation is *justified* or not? Some completely different questions concern the interaction between sentence grammar and discourse grammar. Is sentencehood relevant for the treatment of kataphors? Or would a notion of topic subordination (which applies both below and above sentence level) be sufficient to account for the facts? Furthermore, is it useful to have both the Chomskian Binding Principles, notably Principle C, and the semantic constraints induced by the Properness Condition in the grammar? Is there an overlap, which might be pruned? What are the effects of the semantic constraints below sentence level?

Chapter 6

Presupposition Theory

6.1 Introduction

A discourse theory is a theory of how utterances come to mean what they mean in the mind of a hearer given the conversation in which they take place, and given the hearer's linguistic knowledge and world knowledge. The LDG framework of discourse interpretation embodies this view, even though only part of the parameters that govern discourse interpretation have been fully described and much of the discourse theory remains to be specified. Our approach differs from others in that the emphasis is not so much on the implementation of a given discourse theory, but on the description method itself. LDG can be employed as a specification language for discourse theory in general, and our framework of discourse interpretation may serve as a test-bed for alternative treatments of topics in discourse theory. When we operate out different approaches from their individual formal surroundings and implement them in our framework, it can be seen whether the approaches are compatible and to what extent they overlap.

In this chapter the description method is applied to presupposition theory. In our treatment we combine the approach to presupposition theory proposed in Karttunen (1974), Heim (1983), Beaver (2001) and others, with the approach put forward in Van der Sandt (1992), Van der Sandt and Geurts (1991), and Geurts (1999). The first is often called the *satisfaction* theory of presupposition. It treats presuppositions as tests on the context of interpretation. For a sentence to be felicitous or interpretable, the context of interpretation or 'local context' must entail or satisfy its presuppositions. The local contexts of embedded sentences are established on the basis of the logical structure of the embedding discourse context. Sometimes a presupposition can only be satisfied by strengthening the local context; this is accommodation. The second approach may be called the *anaphoric* theory of presupposition. Its central idea is that presuppositional material must be *bound* in the discourse context of the trigger. If it cannot be bound, it must

be inserted at some semantically accessible location in the discourse context; this is accommodation.

In the satisfaction approach, presupposition satisfaction is defined in terms of entailment, and contexts are modeled as sets of possible worlds or information states (sets of sets of possible worlds). In the anaphoric approach, presupposition resolution is defined in terms of 'binding' or 'matching' of the presupposition in some structured representation of the context, where the semantic language employed is usually DRT or some form of underspecified DRT.[1]

The key notions that we need to formulate the two approaches in our framework of discourse interpretation have been made available in the preceding chapters. *Local contexts* in LDG may serve as 'context of interpretation' in both approaches. On the one hand, they are modeled after Karttunen's concept of local context, which is central in the satisfaction approach. On the other, they are DRSs, representing the semantic information available on the accessibility path from a presupposition trigger to the main DRS. *Satisfaction* can be modeled by means of the entailment relation $\models$, and *binding* by means of properness requirements on discourse meanings and local contexts, in the manner proposed in chapter 5. We will say more about the notion of *accommodation* in a moment.

We will propose a treatment of presupposition theory in LDG that fuses aspects of both approaches. On the one hand, we will follow Karttunen (1974) in requiring that presuppositions must be entailed by the local context of the expression that triggers it. We call this the 'presuppositional condition'. On the other, we recognize that presupposition triggers can be anaphoric elements. In the manner discussed in the preceding chapter, they introduce underspecified discourse markers in their semantic values and in the presuppositional material, and these must find binders in the context of interpretation. Thus, in line with Van der Sandt (1992), presuppositional material introduced by anaphoric triggers must be bound in the local context of the trigger. If it is assumed that presupposition triggers are anaphoric by definition, the essence of the anaphoric approach to presupposition is captured.

The merit of the combined treatment is, very succinctly, that it adds the explanatory power of the satisfaction approach to the concrete predictions of the anaphoric approach. It strenghtens the anaphoric analysis in the sense that presuppositional conditions function as suitability constraints on the resolution of presuppositional elements. It supports a generalisation of the anaphoric account to bridging anaphors. Moreover, it explains what sets binding and accommodation going when it occurs. Presuppositional conditions instigate a reasoning process which, in the case of grammatical discourse, results in their satisfaction in the context of interpretation. Finally, the fusion of the two approaches results in a more complete, and intuitively appealing treatment of accommodation.

[1]See for instance Van der Sandt (1992), Geurts (1995), Keller (1997), Asher and Lascarides (1998b), Blackburn and Bos (1999), and Kamp (2001a).

The basic treatment is laid down in section 6.2. Presupposition theory is a fully integrated part of discourse theory. As before, language users obtain the meaning of an utterance in context by reasoning about discourse descriptions given their grammatical knowledge, their world knowledge, and by employing common sense reasoning and computing preferences over alternative possible readings.

Both the satisfaction approach and the anaphoric account refer to parameters in the discourse theory outside presupposition theory to account for some of their data, be it not the same sets of data. Beaver (1997) points to common sense reasoning and a preference for plausible readings to account for readings in the context of conditional sentences that come out logically stronger than may be expected on the basis of presupposition satisfaction alone. Van der Sandt (1992) introduces constraints on informativity and consistency to account for readings where presuppositional material is bound or accommodated locally, and which otherwise would be too strong. All of these constraints and preferences are already available in the LDG framework for independent reasons, and restrict the number of most preferred possible readings of a given sentence. In section 6.3 we zoom in on the informativity and consistency requirements in particular, and reassess their place in the grammar given their interaction with presupposition and accommodation.

So far, accommodation in our framework was defined as the free specification of world knowledge in B, the implicit background of a conversation. The implicit background models the initial context of interpretation, and what is accommodated in it may very well be distinct from the presuppositional material, as long as it is conducive to the satisfaction of the presupposition or other context-sensitive constraints. This corresponds to *global* accommodation in the satisfaction approach, and to what is called 'context restriction' in Van der Sandt (1988). It differs from the notion of accommodation in the anaphoric approach to presupposition, however. We will continue to refer to it as global accommodation.

In section 6.4 global accommodation is complemented by a form of resource bound accommodation that results in the specification of the semantic content of a discourse, rather than its implicit background. We will call this *content modification*. It corresponds to the concept of accommodation in the anaphoric approach, where presuppositional material is inserted in accessible locations within a discourse representation, outscoping the trigger. Content modification will allow us to model what is generally called 'local' and 'intermediate accommodation' along the lines of the anaphoric approach to presupposition. It also covers 'global accommodation' in the sense of the anaphoric approach.

With the full power of global accommodation and content modification available, the question becomes urgent how the number of possible readings can be restrained. Generally speaking, this will be a matter of adding constraints and specifying preferences. In particular, it will be discussed how some preferences adopted in the anaphoric account, such as the preference for binding over accommodation, and the preference for global over local or intermediate accommodation

can be incorporated in our framework. We will consider how these preferences relate to others in our framework, such as the preference for choosing more plausible readings. In addition, building on Kamp (2001b), we will adopt a preference for minimizing accommodation when it takes the form of content modification.

In section 6.5 the resulting treatment is applied to a series of examples. We focus on cases where through the interaction with independent constraints and preferences in the framework, presupposition triggers are bound 'nonglobally' either by a linguistic antecedent or by accommodated presuppositional material. Broadly speaking, the predictions of our discourse grammar will be those of the anaphoric approach, if the preferences in the grammar are set in accordance with the preferences in Van der Sandt (1992). Presuppositional conditions serve to explain what suitable binders are, and how presupposition interpretation may lead to the specification of general background information. Furthermore, with the help of additional justification conditions on the insertion of background material in specific locations we can account for the abduction of bridging links whenever presuppositional anaphors do not resolve directly to an antecedent present in the discourse context. It may thus be concluded, as we do in section 6.6, that the two approaches are not incompatible and can fruitfully be combined.

6.2 Binding and Satisfaction in Tandem

As announced, the approach to presupposition that we will follow combines aspects of Karttunen's (1974) satisfaction theory of presupposition projection and Van der Sandt's (1992) presupposition projection as anaphora resolution approach. As in Karttunen (1974) *presuppositions* shall be viewed as conditions on local contexts, and the relation of satisfaction shall be implemented as a relation of entailment. Presupposition *triggers* are lexical elements or phrases that come with presuppositional conditions. Suppose a sentence contains a trigger which introduces a certain presuppositional condition. We will say that the sentence *is presupposing* if, when it is uttered as the first sentence in a discourse, the effect of the interpretation process is that the context of interpretation of the sentence entails the presuppositional material. Embedded sentences may be called presupposing if, as a result of the interpretation process, the local context relative to which they are tested entails the presuppositional material.

As before, discourse meanings in our grammar are defined as $B \oplus \sigma_r$, what is taken for granted by the participants in the discourse updated with what is explicitly stated in it. Presupposition satisfaction may induce the partial specification of B, that is, the language user may accept information in B in order to satisfy a given presupposition. This is global accommodation. Information which is globally accommodated ends up in the main DRS of the discourse meaning, and is entailed by it. In what follows, we will use the terms 'local' and 'intermediate accommodation' in line with Van der Sandt (1992), to refer to cases where the

accommodated information ends up in a subordinated DRS and is not entailed by the resulting discourse meaning. A technical treatment of this in LDG is provided in section 6.4.

Discourse interpretation is a reasoning process, and presuppositional conditions are just one of the many constraints that must be satisfied to obtain a number of grammatical readings of a sentence in a given discourse. In satisfying presuppositions, there is interaction with the general constraints in the grammar. Notably, whenever a presupposition trigger is anaphoric in the sense that it introduces a free discourse marker in its semantic value and in the presupposed material, the Properness Condition, introduced in chapter 4, enforces the binding of the marker in the local context of the trigger. The presuppositional condition then is, in effect, a suitability requirement on the resolution of the presuppositional element, and presupposition interpretation results in the selection of a suitable binder. This is a point where our treatment converges with the presupposition projection as anaphora resolution approach. Given a general choice of presupposition triggers as anaphoric elements, and modulo a certain choice of preferences and general acceptability constraints on output readings, the predictions of Van der Sandt (1992) will be obtained, while the concept of presuppositions as conditions on the context of interpretation may be preserved.

Our combined approach will be applied to some simple examples in this section, so that the basic treatment is clear. As a preliminary, we take a closer look at some presupposition triggers in the lexicon (section 6.2.1). After that, we consider a case of plain binding to an antecedent present in the discourse context (section 6.2.2), and some cases where a binder is made available through global accommodation (section 6.2.3). In a brief intermezzo, the interpretation of discourse markers in the universe of the implicit background B, and the distinction between referential and existential readings of definite descriptions is discussed (section 6.2.4). Next, we explain how 'weak' presuppositional conditions, in interaction with the other constraints and preferences of the discourse theory, may result in 'strong' readings (section 6.2.5). We subsequently zoom in on the function of the presuppositional condition as a suitability constraint, illustrating this with a so-called 'partial matching' case (section 6.2.6). Finally, we touch upon the treatment of bridging anaphora (section 6.2.7).

6.2.1 Presupposition Triggers in the Lexicon

We will focus on definite descriptions such as *the man* or *that book*, but the treatment in principle applies equally to other presupposition triggers, such as for example cleft constructions, change-of state verbs, and temporal clauses. We follow Webber (1978) and Heim (1982) in assuming that definites generally must either pick up a discourse referent that is already present in their local context or accommodate such a discourse referent. Definite descriptions, in other words, are much like pronouns in this respect. In the elementary description of the definite

determiner (6.1a) below this is captured by letting their semantic value be an underspecified discourse referent σ_k^π. Properness requirements force σ_k^π to either unify with an existing discourse referent, or to be included in the universe of the background.[2]

Unlike pronouns, definite descriptions also carry descriptive material that is contributed by their common nouns. This material, it is widely assumed, is of a presuppositional character. In line with Karttunen (1974) we require that it is entailed by the local context of the point where it is triggered. In (6.1a) the requirement $\Gamma_{k_1} \models \sigma_{k_2}(\sigma_k^\pi)$ implements this. The result of combining (6.1a) with a description of a common noun is shown in (6.1b). Here the process of taking arbitrary witnesses has resulted in certain instantiations for k, k_1 and k_2.

(6.1) a.

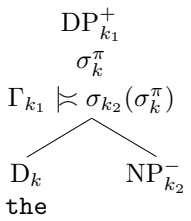

b.

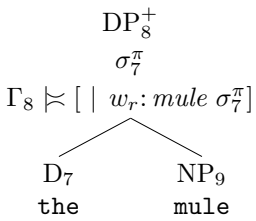

The underspecified, free discourse marker σ_k^π is what makes definites into anaphoric elements. The context-sensitive condition $\Gamma_{k_1} \models \sigma_{k_2}(\sigma_k^\pi)$ classifies them as presupposition triggers. Potential binders of the definite description must 'fit' the description given by the presuppositional condition; the presuppositional condition serves as a *suitability requirement* on the resolution of the definite. Our grammar, in this respect, goes somewhat further than Van der Sandt's (1992) formal treatment, which implements only consistency and informativity as limiting constraints on the matching or insertion of presuppositional material in the context of interpretation.

Note that the underspecified discourse marker in (6.1a) and (6.1b) is like a pivot, linking the semantic content of the definite to the presupposed material. The presupposition can only be satisfied given some binding of the shared discourse marker in the common local context, and the properness requirements of our grammar enforce binding. Thus, as in the anaphoric approach to presupposition projection, problems potentially arising from stating presuppositions of definite descriptions as existentially quantified statements *independently* of the asserted material do not arise. Neither do we predict the projection of certain unwanted universally quantified presuppositions, as occuring in the approach of for example Heim (1983).[3]

[2] Alternatively, as in Van Leusen and Muskens (2003), binding requirements of the form $[\sigma_k \mid] \sqsubseteq \Gamma_{k_1}$ can be incorporated in the elementary description of the definite determiner.

[3] Heim (1983) predicts that the sentence in (6.2a) presupposes or projects (6.2b). This basically results from having a free variable (or unbound discourse marker) in the presuppositional material triggered by 'his bicycle'.

The lexemes and linguistic constructs that are usually called presupposition triggers are known to be quite a diverse lot.[4] Given the tools available in our grammar, they may each be classified in the lexicon according to their specific characteristics as context-sensitive elements. Thus, though anaphoricity may well be a general characteristic of presupposition triggers, the extent to which their interpretation in context exhibits binding effects depends on their precise semantic characterisation, and this may be varied. To illustrate this, (6.3) shows a possible lexical description of the change-of-state verb 'sell', which is commonly taken to presuppose that the agent of the selling event is in a state of owning the thing sold. As in chapter 4, we introduce some notation that suggests an event-semantic analysis of the lexeme: e_k, e_{k_1} are discourse markers of type eventuality, which covers both states and events, e_{k_1} is the main eventuality over which the semantic value of k_1 is predicated, and $e_k \oslash e_{k_1}$ means that (pre-state) e_k obtains at the moment when (event) e_{k_1} starts.[5]

(6.3)

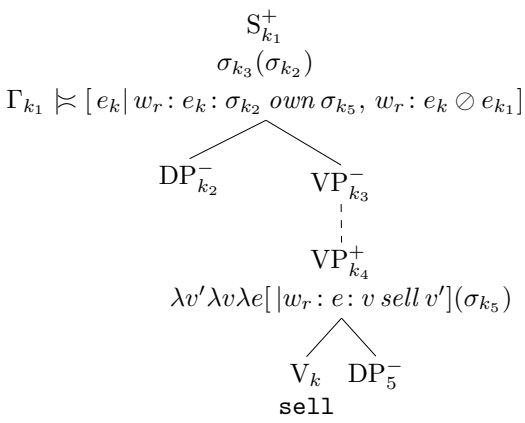

With verbal triggers such as these, it seems less clear that they are directly anaphoric, like definite descriptions. The presuppositional condition

$$\Gamma_{k_1} \mid\!\!\sim [\, e_k |\, w_r : e_k : \sigma_{k_2}\ own\ \sigma_{k_5},\ w_r : e_k \oslash e_{k_1}]$$

reflects this. The discourse marker e_k, representing the state of ownership, is fully specified and bound within the presuppositional material. The condition is

(6.2) a. *A fat man was pushing his bicycle.*

b. *Every fat man owns a bicycle.*

Heim considers this an unintuitively strong prediction, and points out the desirability of having an alternative treatment in which accommodation results in the more acceptable reading that there was a fat man who had a bicycle and was pushing it. This in fact corresponds to the most preferred possible reading of (6.2a) predicted by our treatment. It results from resolving the possessive pronoun to 'a fat man' and globally accommodating that he owns a bicycle. Our predictions are in line with those of Van der Sandt's anaphoric account here.

[4] See for instance Levinson (1983), chapter 4, for an overview.

[5] This is very sketchy and only serves to illustrate the point; we are ignoring complications potentially arising from the interaction with tense, VP-modifiers, and quantification in an event-semantic analysis here.

satisfied when the existence of a state of this kind can be inferred from the local context. Binding effects are predicted to result only indirectly, from the binding of σ_2 and σ_5 in the discourse context. Further research would have to show in how far this analysis fits the observable facts.

There are also cases in which the presuppositional condition seems to be superfluous and the presupposition trigger is essentially an anaphor. The temporal adverb 'then' is an example of this. Following Webber et al. (1999) and Webber et al. (2003) the discourse adverbial 'then', as occurring in, for instance, "John handed in his resignation and went home. Then he made himself a stiff drink." may be taken to presuppose a telic eventuality, and to assert that the eventuality described by the argument clause starts after its culmination.

This may be expressed by means of the following lexical description, where e_{k_2} represents the eventuality described by the argument clause (S_{k_2}) and σ^{π}_{k} is an underspecified discourse marker which stands for the presupposed telic eventuality.[6] The temporal relation between the presupposed eventuality and the asserted one is expressed by $w_r{:}\,\sigma^{\pi}_{k} <_t e_{k_2}$. We assume that sortal restrictions on σ^{π}_{k} can be added, to make it into a discourse referent that can only be occupied by telic eventualities. The main eventuality of the whole sentence is represented by e_{k_1}, and the condition $w_r{:}\, e_{k_2} \subseteq e_{k_1}$ adds some referential fine structure. We assume e_{k_2} is bound in σ_{k_2}.

(6.4)

$$\begin{array}{c} \mathrm{S}^{+}_{k_1} \\ [\, e_{k_1} \,|\,] \oplus \sigma_{k_2} \oplus [\;|\; w_r{:}\, e_{k_2} \subseteq e_{k_1},\; w_r{:}\, \sigma^{\pi}_{k} <_t e_{k_2}\,] \\ \diagup \qquad\qquad \diagdown \\ \mathrm{Adv}_k \qquad\quad \mathrm{S}^{-}_{k_2} \\ \texttt{then} \end{array}$$

As with pronouns, the underspecified discourse marker σ^{π}_{k}, which is free in σ_{k_1}, must find a binder in Γ_{k_1}. As we saw in the preceding chapter, the Properness requirement says that discourse meanings must be proper. It follows from this that σ^{π}_{k} must be bound in the discourse meaning. Given the composition of local contexts, the embedding context of σ_{k_1} in the discourse meaning is precisely the local context of the discourse unit, Γ_{k_1}. Hence σ^{π}_{k} must be bound in Γ_{k_1}. The presupposition trigger is basically a temporal anaphor. Sortal properties constrain its resolution.

As a final example, we would like to mention triggers such as 'too', which are also known as *focusing adverbs*. It might be argued that in these cases, focus interpretation overlaps with presupposition theory, and the 'focus frame' or 'background' that can be inferred from the focus articulation of the argument of the adverb constitutes the presuppositional material. It is commonly assumed

[6] Remember there is a family of functions σ^{α}, assigning semantic values of any given type α to a node in a tree structure. We employ this multiplicity here by letting σ^{π} generate an underspecified discourse marker at node k, while $\sigma^{\tau\tau}$ produces the 'standard' or 'main' semantic value of the adverbial.

that focus frames must in some sense be 'given' in the context of interpretation.[7] One way of implementing this could be to treat them just like presuppositions, and require that focus frames be entailed by the local context of the clause that triggers them.

To illustrate this, consider the following lexical description of 'too'. We use fv_{k_2} to denote the focus value of the argument of the adverb. A focus frame may in principle be obtained from this by replacing the focus value in the semantic content of the clause with a free variable x_k of the same type α.[8] The presuppositional condition can then be stated as $\Gamma_{k_1} \models \{x^\alpha_k / fv^\alpha_{k_2}\}\sigma_{k_2}$, where $\{x^\alpha_k / fv^\alpha_{k_2}\}\sigma_{k_2}$ indicates the substitution of $fv^\alpha_{k_2}$ by x^α_k in σ_{k_2}. Satisfying the condition should result in finding some specific value that matches x^α_k. If $\alpha = \pi$, x^α_k is an underspecified discourse referent and the Properness requirement applies to it as usual.

(6.5)

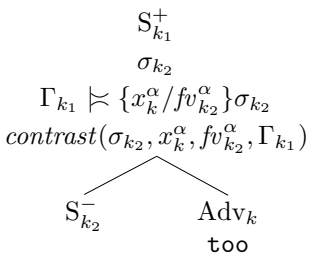

Besides givenness of the focus frame, focus interpretation is usually taken to involve another requirement, namely distinctness or contrastivity of the focus value with the value that resolves x_k.[9] We would like to suggest that the semantic relation of contrast, whatever its precise definition is, is a context-dependent notion. The schematic representation of the condition $contrast(\sigma_{k_2}, x^\alpha_k, fv^\alpha_{k_2}, \Gamma_{k_1})$ reflects this. In fact, the givenness requirement on the focus frame could be viewed as a component or corollary of the contrastivity requirement. The focusing adverb exhibits binding effects to the extent that $\{x^\alpha_k / fv^\alpha_{k_2}\}\sigma_{k_2}$ contains free discourse markers. As for its accommodation potential, we can say very little; as with pronouns and definite descriptions, it depends on the interaction with other constraints associated with the trigger.

Clearly, much of the above is tentative and sketchy. The behaviour of various classes of presupposition triggers is the topic of extensive research, the discussion of which would take too much space in this chapter. Furthermore, a full treatment in our grammar would require at least a definitive specification of a fragment of English in event semantic compositional DRT, and an implementation of focus theory. Both will have to wait for some later time. We hope to have shown,

[7] Starting from Jackendoff (1972) and Dahl (1974). See also Rooth (1985), Krifka (1993), Partee (1991), Geurts and van der Sandt (1997).

[8] Compare Gardent and Kohlhase (1995), Pulman (1997) in the context of higher order unification theory and matching, and see Kamp (2004), Riester (2005), for the construction of focus frames in DRT.

[9] Compare Kaplan (1984), Rooth (1992).

however, that with the means standardly available in our discourse semantics, such as context-sensitive constraints and properness requirements, a unified and declarative treatment of linguistic presupposition can be provided, while doing justice to the diversity of the triggers in the lexicon.

6.2.2 Reasoning about Presuppositional Conditions

After this brief look at the lexicon, let us zoom in on the interplay between binding, presupposition satisfaction and accommodation in the interpretation process. This can best be illustrated with some simple examples. Consider the first one of these, given in (6.6).

(6.6) *Pedro has a mule. The mule is happy.*

As before, upon hearing each sentence the language user constructs an input description, and increments the existing discourse description with it. The discourse description representing the hearer's knowledge of the discourse just after processing (6.6) can be represented as in (6.7). For transparancy, we assume some node identifications have already taken place, and the syntax and semantics of the predicate 'is happy' is simplified.

(6.7) Δ_2:

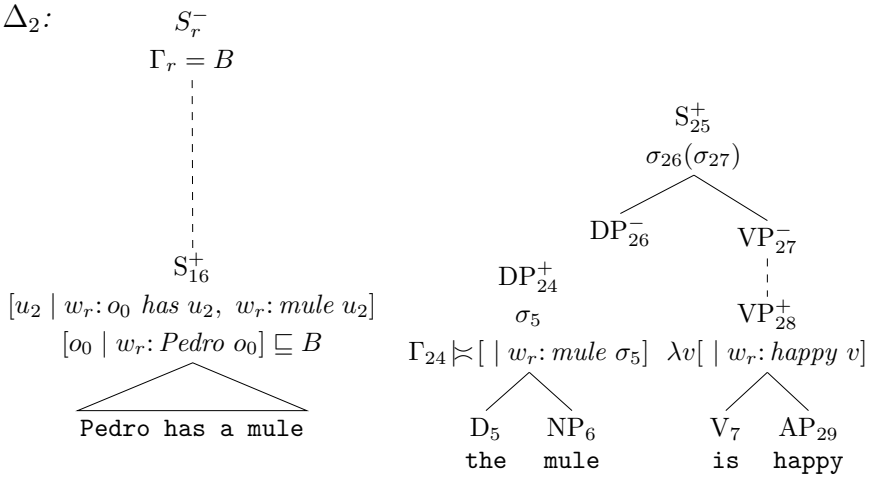

Reasoning on the basis of his grammatical knowledge, world knowledge, and preferences the hearer may further parse the input sentence, and select a suitable discourse relation, expressing its link to the preceding discourse. Discourse relation *elaboration* seems an appropriate candidate. On the assumption that the discourse is temporarily ended, he may subsequently infer the discourse description shown in (6.8).

(6.8) Δ_2^{rcl}:

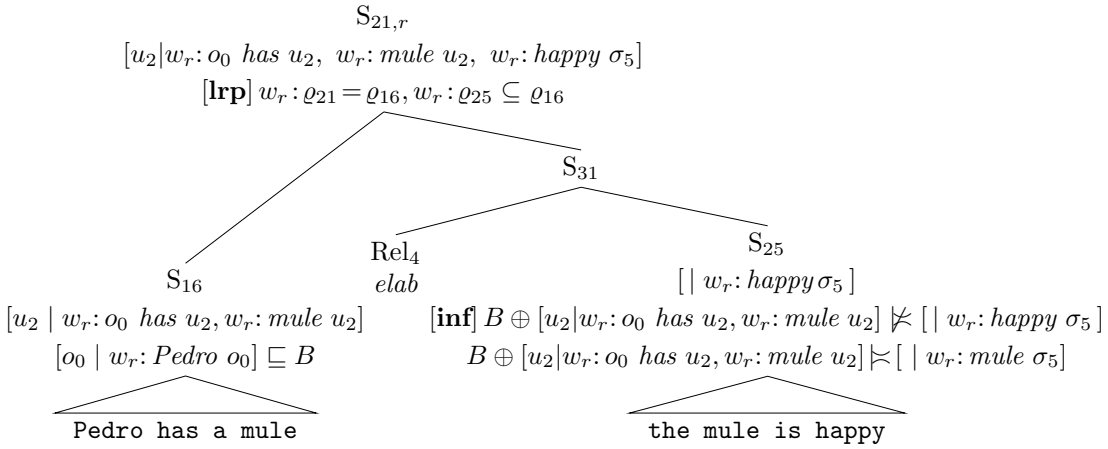

As before, our hearer may obtain the meaning of the discourse by reasoning about $B \oplus \sigma_r$ given satisfaction of the constraints triggered by the various context-sensitive elements in the discourse. The discourse meaning that can be inferred from (6.8) is the underspecified DRS (6.9a), where the conditions in (6.9b-d) hold. Condition (6.9b) was introduced by the name 'Pedro', (6.9c) is the local informativity condition introduced by the elaboration relation, and (6.9d) is the presuppositional condition triggered by the definite description.

(6.9) a. $B \oplus [u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon \mathit{happy}\ \sigma_5]$
b. $[o_0 \mid w_r\colon \mathit{Pedro}\ o_0] \sqsubseteq B$
c. $B \oplus [u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2] \not\models [\,\mid w_r\colon \mathit{happy}\ \sigma_5]$
d. $B \oplus [u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2] \models [\,\mid w_r\colon \mathit{mule}\ \sigma_5]$

As discussed in chapter 4, the requirement in (6.9b) induces global accommodation of the presence of an individual named 'Pedro'. Furthermore, the informativity requirement in (6.9c) prevents the accommodation of information that supports that σ_5 is happy in the local context of the second sentence, for any possible binding of σ_5. Now what about the presuppositional condition (6.9d)? Given (6.9b), axiom $\mathcal{A}15$, and idempotency of $\oplus$, (6.9a) can be strengthened to

$$B \oplus [w_r\ o_0\ u_2 \mid w_r\colon \mathit{Pedro}\ o_0, w_r\colon o_0\ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2,\ w_r\colon \mathit{happy}\ \sigma_5\,]$$

Given the properness requirement $\mathcal{A}18$, the hearer may infer from this that

$$\mathit{proper}(B \oplus [w_r\ o_0\ u_2 \mid w_r\colon \mathit{Pedro}\ o_0, w_r\colon o_0\ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2,\ w_r\colon \mathit{happy}\ \sigma_5\,])$$

With respect to the underspecified discourse marker σ_5, he may subsequently infer

$$[\sigma_5 \mid\ ] \sqsubseteq B \oplus [w_r\ o_0\ u_2 \mid w_r\colon \mathit{Pedro}\ o_0, w_r\colon o_0\ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2,\ w_r\colon \mathit{happy}\ \sigma_5\,]$$

and conclude that $\sigma_5 \in \{w_r, o_0, u_2\}$, or that σ_5 is in the universe of B but not in $\{w_r, o_0, u_2\}$. Of the four potential binders, w_r can immediately be excluded on the basis of sortal constraints. Of the other binding options, identifying σ_5 with o_0 or assuming $\sigma_5 \notin \{w_r, o_0, u_2\}$ results in dispreferred readings of the sentence. The most preferred reading results from identifying σ_5 with u_2; the definite description then resolves to 'the mule Pedro owns'. It is clear that for this choice the presuppositional condition in (6.9d) holds:

$$B \oplus [u_2 \mid w_r\colon o_0 \ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2] \models [\mid w_r\colon \mathit{mule}\ u_2]$$

The corresponding discourse meaning (6.10a) is constrained as in (6.10b).

(6.10) a. $B \oplus [w_r\ u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon \mathit{happy}\ u_2]$

b. $[o_0 \mid w_r\colon \mathit{Pedro}\ o_0] \sqsubseteq B \wedge$
$B \oplus [u_2 \mid w_r\colon o_0\ \mathit{has}\ u_2, w_r\colon \mathit{mule}\ u_2] \not\models [\mid w_r\colon \mathit{happy}\ u_2]$

This is a reading in which the presupposition is 'caught' or 'bound' by previous linguistic material. Nothing needs to be accommodated in B in order to satisfy the presuppositional condition. The discourse in (6.6) can be assumed to have been uttered in an implicit context that does not contain any reference to mules at all.

As we saw, there are two other possibilities to bind σ_5 and satisfy (6.9d), resulting in possible but dispreferred readings of the discourse. One of these is to identify σ_5 with o_0 and globally accommodate that Pedro is a mule. It might be argued that lack of sufficient justification for what must be accommodated explains why this reading is dispreferred. Only in specific (story-telling) contexts, mules may own one another. Presumably, the selection of such a context in the absence of any linguistic or nonlinguistic evidence supporting it is not very well justified, and dispreferred when better alternatives are available. Another reason why this reading is dispreferred might be that it implies that there is more than one mule present in the context of interpretation, possibly within the relevant domain of topical elements, while the descriptive material in 'the mule' is not distinctive enough to select one of these uniquely. It could be argued that the use of the definite description in that situation is uncooperative on the part of the speaker, because the hearer is not given enough information to disambiguate. On the hypothesis that the speaker does act cooperatively, these readings cannot be what he intends to convey by his utterance, hence they are dispreferred.

The other dispreferred reading results when σ_5 is taken to refer to an individual in the implicit background other than Pedro, and it is globally accommodated that this individual is a mule. The fact that the descriptive material is not distinctive enough to select a candidate uniquely may explain the dispreferredness of this reading as well. On top of this, given the specific interpretation of discourse-markers in the universe of the implicit background B in our semantics, the definite

description would receive a deictic or referential interpretation—a point somewhat further discussed below. As with pronouns, it might be argued that deictic or referential readings require specific evidence from the linguistic or nonlinguistic context of interpretation; in its absence accommodation is not suffiently justified.

Generally speaking, the hearer's *preferences* govern the interpretation of these and similar cases, and should explain why some readings are preferred and others are dispreferred. A direct explanation of the judgements in the example at hand is obtained if we follow the anaphoric account to presupposition of Van der Sandt (1992) and adopt an explicit preference for 'binding' presuppositional material over 'accommodating' it. In our treatment this corresponds to adopting a preference for binding a presupposition to a discourse marker introduced directly by a linguistic antecedent over binding it to an accommodated one. This produces the correct predictions for discourse (6.6): the most preferred reading is the one in which the presupposition is satisfied and bound relative to 'a mule'. In the dispreferred readings it is satisfied and bound relative to accommodated material.

The preference for binding to a linguistically introduced binder may itself be explained in terms of a preference for more cohesive interpretation, on the assumption that resolution to explicit antecedents is generally perceived to be more cohesive than resolution to inferred information. We leave open whether this is so. As we have tried to indicate in the above discussion, it is also possible to explain the judgements in terms of other preferences, such as the preference for selecting most plausible or uncontroversial background information, and a preference for taking the speaker to avoid ambiguous reference. All of this may be seen as the starting point for further research, where the question what makes one reading more costly than the other is central. A theory of justification, or justified acceptance of information in B is part of this.[10] In section 6.4.3 we will take up the issue of the interaction between the various preferences in our analysis again.

6.2.3 Binders can be Globally Accommodated

In the previous example, the most preferred reading was one in which the definite description was bound by an antecedent present in the preceding discourse. This is not necessarily so, however. A case where global accommodation of the binder and the presuppositional material is clearly preferred is given in (6.11a). We assume the hearer constructs a discourse description for this text as before. Local informativity requirements are not relevant for the present discussion and we

[10]In for instance Dalrymple et al. (1994), Zeevat (1999), and Blutner (2000) the idea that there is an overall preference for stronger possible readings over weaker ones is considered. The three readings discussed here, however, are not ordered in terms of logical strength relative to each other, so we cannot use them as a test case for this hypothesis. We will come back to this when discussing local accommodation further on.

ignore them for the time being. Given satisfaction of the condition triggered by the proper name in (6.11a), the presuppositional constraint triggered by 'the sun' may be given as in (6.11b). A consequence of the Properness requirement is shown in (6.11c).

(6.11) a. *Pedro is asleep. The sun is shining.*
b. $B \oplus [\, w_r\, o_0 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon \mathit{asleep}\ o_0\ ] \mid\!\sim [\ \mid w_r\colon \mathit{sun}\, \sigma_4]$
c. $\mathit{proper}(B \oplus [\, w_r\, o_0 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon \mathit{asleep}\ o_0,\ w_r\colon \mathit{shining}\, \sigma_4])$

Given (6.11c) it follows that σ_4 is bound in the universe of the discourse meaning. For reasons similar to those mentioned before, identifying σ_4 with o_0 and accommodating that Pedro is the sun results in a dispreferred reading. The only preferred option is to assume the discourse referent, σ_4, to be already present in the background B, in other words, to accommodate $[\sigma_4|\ ] \sqsubseteq B$. In addition, the background must contain sufficient material to satisfy the presupposition in (6.11b). A simple way to satisfy this requirement is if in fact $[\ \mid\ w_r\colon \mathit{sun}\ \sigma_4] \sqsubseteq B$. Thus, the hearer may infer that discourse meaning (6.12a) is constrained as in (6.12b).

(6.12) a. $B \oplus [\ w_r \mid w_r\colon \mathit{asleep}\ o_0,\ w_r\colon \mathit{shining}\, \sigma_4]$
b. $[\ o_0 \mid w_r\colon \mathit{Pedro}\ o_0\ ] \sqsubseteq B \wedge [\ \sigma_4 \mid w_r\colon \mathit{sun}\ \sigma_4\ ] \sqsubseteq B$

In the preceding chapter (section 4.4.7) it was assumed that all discourse referents in the universe of the implicit background of a conversation are interpreted *referentially*. It was assumed that discourse participants interpret the discourse with respect to some state i_0 which they take the discourse to be about, and the discourse markers in the universe of B receive their interpretation from i_0 directly. The state i_0 functions much as an *external anchor* in the sense of Kamp and Reyle (1993), pp. 246–248. Given that treatment, the definite description in (6.11a), whose binder is globally accommodated, is interpreted referentially. In (6.11a) this is intuitively appropriate: the speaker, when he utters 'the sun is shining', is referring to a particular entity present in the world he is talking about.

What justifies global accommodation of a binder can be more than just common sense knowledge or knowledge about the world; it can also be information that is provided explicitly in the discourse. Consider the conditional sentence in (6.13a), which in one of its possible readings presupposes that John has children. We agree with Van der Sandt (1992) that this is its most preferred reading. The fact that in continuations such as (6.13b) the pronoun 'they' can be used to refer to John's children attests that the presupposition is globally accommodated; if it were locally accommodated or satisfied in the scope of the conditional it could not bind the pronoun.

(6.13) a. *If John has grandchildren, his children must be happy.*

Van der Sandt (1992), ex.(27), p.349

b. *They always wanted off-spring.*

Our treatment accounts for this reading as follows. The definite description 'his children' introduces a presuppositional condition on its local context. Given resolution of the possesive pronoun 'his' to 'John' that presupposition can be paraphrased as 'there are individuals who are Johns children'. The local context of the triggers, established by the antecedent of the conditional, includes the information that John has grandchildren. The information that individuals who have grandchildren have one or more children too may readily be accepted in B, being common knowledge. It then follows that John has one or more children. We assume that the fact that this information can be inferred from common knowledge plus the information provided by the embedding discourse is sufficient to justify the accommodation of a suitable binder in the local context of the presupposition trigger.

As the continuation in (6.13) shows, the accommodated material must be accepted in the implicit background in B. If it were not, it could not bind the anaphor 'they'. The inference that John has children builds crucially on information from the antecedent of the conditional which is not globally available, however. So one might ask what justifies *global* accommodation of the binder in the reading discussed.

For one, it may be assumed that a general preference for global accommodation is at work. It is a widespread assumption that a preference of this kind is part of a language user's preference system, see for instance Heim (1983) and Van der Sandt (1992). This means that, unless grammatical constraints or other preferences interfere, a reading resulting from global accommodation is predicted to be the most preferred. Another option is to try and explain the effect in terms of constraints or preferences already available to the hearer for independent reasons, such as the preference for more plausible readings, and informativity or appropriateness conditions associated with conditional constructions. We would suggest that in the example at hand, the information that is accommodated (John has children) is a natural precondition of what is stated in the antecedent of the conditional (John has grandchildren). If one considers the possibility of John having grandchildren, one normally adopts a background in which he has children. Thus, there is a sense in which the global reading corresponds to an optimally informative or appropriate use of the conditional.

Since our discourse semantics in its current form allows only global accommodation, our theory predicts only the reading discussed here. Given an extension of the semantics that allows for local accommodation, to be provided in section 6.4, the grammar predicts possible readings resulting from local accommodation for (6.13b) as well. The continuation in (6.13b), however, enforces the selection of the reading in which it is globally accommodated that John has children as the

only possible one.

6.2.4 Referential and Existential Readings of Definites

While for the purpose of this chapter there is no particular need to change the interpretation of discourse markers in the universe of B that we adopted, we would like to point out that the semantics allows for variations, and a set-up may be chosen that results in a more diversified interpretation of nominal elements whose binders are in the implicit background of the the conversation. In particular, a treatment might be chosen in which referential interpretations are not assigned *generally* to definite descriptions whose binder is globally accommodated. Consider the following example.

(6.14) *A girl fainted. Her boy-friend called the doctor.*

It may be claimed that in addition to a referential reading, an existential reading of 'the doctor' in (6.14) is available. In that reading the speaker is not talking about a particular individual, but about any given individual that fits the description. He just conveys that a doctor was called.[11] In order to account for this, a treatment could be adopted in which some discourse referents declared in the universe of the global background are referential while others are interpreted existentially. This can be achieved by employing the existing difference between o-markers and u-markers to distinguish between referential and existential interpretations, rather than 'background' and 'new' markers. The following semantic axiom may be introduced to regulate this; '*ref*' is a πt-type predicate which marks a discourse referent as referential.

$\mathcal{A}22\ \ \forall k.ref(o_k) \wedge \forall k.\neg ref(u_k)$

To allow existential readings of background material, u-markers must no longer be excluded from the universe of the global background of the conversation. Arguably, it may also be required that o-markers are declared in the universe of the global background only.[12] Thus, instead of

[11] There is a convergence here with Donnellan's (1966) distinction between a referential, and an attributive use of definite descriptions. With Donnellan's *referential use*, the definite description serves to (let the audience) pick a given, particular object or individual. The descriptive material is presumed to apply to that object or individual, but is not essential in establishing its identity. With the *attributive use*, the attributed property is essential. It is assumed that there is an individual to which the description applies. A definite description of the form 'The F' might be paraphrased as "the F, whichever or whoever it may be,...". While Donnellan's 'referential use' seems to fit our treatment of o-markers as externally anchored referents, his 'attributive use' seems to correspond best to our 'existential readings', where a u-marker must be accommodated.

[12] Though in fact, it is not necessary to state this explictly. In the fragment of our grammar *names* are the only elements introducing o-markers, and the the declaration of these markers in the universe of B is achieved by assigning names a context-sensitive condition which requires the discourse marker and its description to be in the universe of B.

$\mathcal{A}19$ $\mathit{proper}(B) \wedge \neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$

We would now have

$\mathcal{A}19''$ $\mathit{proper}(B) \wedge \forall k\, \mathsf{fst}(B)(o_k) \wedge \mathsf{fst}(B)(w_r)$

In addition, the truth definition we had ('a DRS K is said to be *true in state i with respect to background B* if $\mathit{true}(\langle \mathsf{fst}(K) - \mathsf{fst}(B), \mathsf{snd}(K)\rangle)(i)$ holds.') may be replaced with one that excludes the set of all o-markers from interpretation as an existentially bound variable: 'a DRS K is said to be *true in state i with respect to a set of externally anchored referents 'ref'* if $\mathit{true}(\langle \mathsf{fst}(K) - \mathit{ref},\ \mathsf{snd}(K)\rangle)(i)$ holds.' As before, it would be assumed that discourse participants interpret the discourse with respect to some state i_0 which they take the discourse to be about, and o-markers receive their interpretation from i_0 directly.

Whether the elements of a certain word class select referential or existential readings, or allow both, is regulated in the lexicon of the grammar. For instance, in our fragment names carry o-markers, and indefinites introduce u-markers. The first will consequently be interpreted referentially, the second existentially.[13] Elements which introduce underspecified discourse markers, such as definite descriptions and personal pronouns, allow for both referential and existential interpretations. They can select either o-markers or u-markers as a binder. Elements which get deictic readings only, such as 'I', 'you', and now' could be assumed to carry o-markers. Clearly, there is a vast diversity and subtlety in the interpretation of nominal expressions. We cannot hope to address this in the scope of this chapter. We merely want to show that a certain measure of fine-tuning can be provided in the semantic language we employ.[14]

6.2.5 Accommodating Weak and Strong Readings

In our grammar, global accommodation is the specification or abduction of general background information. What is globally accommodated can be any information that is sufficient to satisfy the presupposition, and which the language user is willing to accept as background information, given his world knowledge and preferences. It need not be precisely the presuppositional material. Crucially, what is globally accommodated can be logically *stronger* than what is minimally required to satisfy a given presuppositional condition. Any specification of the

[13] It is well-known that indefinites can be used referentially, as in 'John invited a certain girl', where 'a certain girl' refers to a specific individual known to the speaker and hearer. The expression is then called a *specific* indefinite. This reading can be obtained by allowing the identification of the u-marker introduced by the indefinite with an accessible o-marker (under condition of sufficient justification). The indefinite will 'inherit' a referential interpretation from the o-marker. The identification with an accessible o-marker is not enforced by any constraint and will typically only occur when it is justified on the basis of available background knowledge.

[14] Compare Maier (2006) for a similarly sophisticated treatment of deixis and referential and descriptive uses of names in Layered DRT.

background that satisfies the condition will do, in this respect. What specifications count as most preferred choices is determined on the basis of *all* constraints in the hearer's grammar and the total of his preferences; it does not depend on the presuppositional condition alone. The interpretation process as a whole may thus result in the selection of a reading that is stronger than the presuppositional condition requires.

That the material that is in fact accommodated in our approach is not always the *weakest* DRS such that the entailment requirement is met with is shown in (6.15).

(6.15) a. *If the weather turns out to be good, the king will be overjoyed.*
b. $B \models [\ |\ w_r\colon \mathit{weather}\ \sigma_2]$
c. $B \oplus [\ |\ w_r\colon \mathit{good}\ \sigma_2\ ] \models [\ |\ w_r\colon \mathit{king}\ \sigma_9]$
d. $B \models [\ |\ [\ |\ w_r\colon \mathit{good}\ \sigma_2\ ] \Rightarrow [\ |\ w_r\colon \mathit{king}\ \sigma_9]\ ]$

The definite description 'the weather' in (6.15a) evokes the presupposition (6.15b); 'the king' evokes the one in (6.15c). We concentrate on the latter. Notice that (6.15c) is equivalent to (6.15d). Thus, in order to satisfy the presupposition triggered by 'the king', it would be sufficient to assume that the background contains σ_9, plus a condition that can be paraphrased as 'if the weather turns out good σ_9 is king'.[15] But clearly this is not what is accommodated in ordinary contexts, where plain 'σ_9 is king' is preferred. In the first case we would have (6.16a), in the second (6.16b)—given satisfaction of (6.15b).

(6.16) a. $[\,\sigma_9 \mid [\mid w_r\colon \mathit{good}\ \sigma_2\ ] \Rightarrow [\ \mid\ w_r\colon \mathit{king}\ \sigma_9]\,] \sqsubseteq B$
b. $[\,\sigma_9 \mid\ w_r\colon \mathit{king}\ \sigma_9] \sqsubseteq B$

We agree with Beaver (2001) that accommodation, while constrained by requirements that derive from properties of the linguistic system, is also a matter of common sense reasoning. The linguistic requirement for a felicitous utterance of (6.15a) is that B entails 'if the weather turns out good σ_9 is king', but there are many ways in which this requirement can be met. Common sense reasoning is needed to pick out the most likely of these ways, or in other words, to develop preferences over potential readings. In the case of conditional statements such as the one accommodated in (6.16a) it should help to gauge the meaningfulness or naturalness of the dependency expressed in them.[16]

[15] This prediction is dubbed 'the proviso problem' in Geurts (1995) and Geurts (1998). It arises when a presupposition χ is triggered relative to a local context of the form $c[\varphi_1][\varphi_2]\ldots[\varphi_n]$, $n \geq 0$, where c is the global context and each $[\varphi_i]$ an update. In this situation only a presupposition of the form $(\varphi_1 \wedge \varphi_2 \wedge ...\varphi_n) \rightarrow \chi$ will survive, whereas the intuitive inference will often be that χ is true.

[16] Interesting matter for comparison is the treatment of presupposition resolution in conditional sentences in Asher and Lascarides (2003). They argue that a *general* preference for

While the specification of the hearer's preference system is an open research task, and we have nothing to contribute to the theory of commonsense reasoning, it is clear that by integrating presupposition theory in a full-fledged discourse theory, the occurrence of logically stronger readings than is necessary to satisfy a given presuppositional condition can straightforwardly be accounted for. It is the effect of interaction with independent constraints and preferences in the discourse theory.[17] This reconciles our Karttunen (1974)-style treatment of presupposition with the predictions of the anaphoric account.

6.2.6 Partial Matching and Suitability

In the preceding we encountered some examples in which the material associated with a definite description could either be unified completely with material already in the linguistically generated part of the local context of that description or had to be accommodated in its entirety. For good measure we also give an example where some material (the discourse referent) is unified with existing material, while other material must be presumed to be in the background. These are called *partial matching* cases in the anaphoric approach to presupposition.[18] Consider (6.17a).

(6.17) a. *If John marries a woman, the unlucky female will be unhappy.*

b. $B \oplus [w_r\, o_1\, u_3 \mid w_r{:}\, \mathit{John}\ o_1,\ w_r{:}\, o_1\ \mathit{marries}\ u_3,\ w_r{:}\, \mathit{woman}\ u_3]$

c. $[\ \mid w_r{:}\, \mathit{female}\ u_3,\ w_r{:}\, \mathit{unlucky}\ u_3]$

d. $[u_9 \mid w_r{:}\, \mathit{woman}\ u_9] \Rightarrow [\ \mid w_r{:}\, \mathit{female}\ u_9]$

e. $[u_9 \mid w_r{:}\, o_1\ \mathit{marries}\ u_9] \Rightarrow [\ \mid w_r{:}\, \mathit{unlucky}\ u_9]$

The local context of *the unlucky female* in this sentence is given in (6.17b) and it is clear that its discourse referent can be taken to be equal with u_3. If this identification is made, (6.17c) must be entailed by (6.17b) and the common sense

more plausible or likely interpretations, as proposed in Beaver (1996), is not enough to account for all of the relevant data, and that rhetorical relations (among which are various types of consequence relations) determine what we should be measuring the plausibility of. Rephrasing somewhat freely, conditional constructions come with their own specific constraints on what is plausible, and this overrules the general notion of plausibility.

[17] This point is made in Asher and Lascarides (1998b) and Asher and Lascarides (2003) as well, although they do not formulate presupposition theory in terms of entailment conditions on local contexts. Rather, they build on the anaphoric approach to presupposition and provide a powerful generalisation over Van der Sandt's 'binding' and 'accommodation' cases in SDRT. The overall interaction with discourse theory explains that some readings are more coherent than others, hence more preferred.

[18] Partial matching is extensively discussed in Van Deemter and Krahmer (1998). They provide a generalised treatment of NPs-as-anaphors, which covers, among others, subsectional anaphora and bridging anaphora. The latter differ from the cases discussed in this section, in the sense that a *binder* must be also be accommodated or inferred from existing material.

reasoning component must abduce one or more plausible conditions such that, if these conditions are assumed to be in B, the desired entailment holds. Clearly, one condition that can easily be assumed to be in B is (6.17d), 'women are females', and a more contentful assumption is (6.17e), 'whoever is married by John is unlucky'.[19] If these assumptions are made, the required entailment holds.

Examples like these illustrate in particular that with anaphoric triggers, the presuppositional condition serves as a *suitability requirement* on the resolution of the anaphor. A suitable binder is a discourse marker whose occupant can be inferred to fit the description given the context of interpretation. The suitability test is what sets the abduction of background information going. It is a fully integrated part of the semantics of presupposition triggers. This can be seen as an enrichment of the anaphoric approach to presupposition of Van der Sandt (1992), where a suitability requirement of this kind is not implemented.[20] The selection of suitable antecedents there results from interaction with informativity and consistency requirements, and discourse properties like salience and the contribution of nonlinguistic knowledge. Thus, to some extent the notion is covered in the pragmatics of the discourse theory surrounding DRT, but this is not formally spelled out.

6.2.7 Bridging

One of the things the preceding discussion shows is that in our reasoning model, general background knowledge can freely be abduced or accommodated. If the presence of binders can subsequently be inferred on the basis of what is asserted and accepted as background information, accommodation of these binders is justified in principle. This is attractive in that it offers a natural account of *bridging* anaphora, a phenomenon that pervades natural language.[21] Bridging anaphors are expressions whose binder is not introduced in the universe of the context of interpretation through an antecedent in the preceding discourse, but which can be assumed to be present on the basis of what we know about the elements that

[19] We require that u_9 is a fresh witness for the discourse referent in (6.17d) and (6.17e). Notice that u_9 is not generated at a node named 9; in fact it must be avoided that there is a node named 9 in the discourse structure.

[20] Though compare Van Deemter and Krahmer (1998), who present a modified version of Van der Sandt (1992) and integrate a semantic suitability test in their resolution algorithm. In their treatment, definite descriptions introduce not only an ordinary discourse marker, but also a fresh *context set*, which in the resolution procedure is equated with a previously introduced discourse marker, or with the entire domain of discourse. Thus, the context set identifies the antecedent of the definite description. Each potential antecedent comes with a so-called *value-set*, which may be circumscribed as a set of linguistically generated chacterising properties. The presuppositional material of the definite description is checked against that set; (generalised) binding in principle succeeds if there is a non-empty intersection.

[21] See for instance Clark (1977), Hobbs (1979), Charniak (1983), Van Deemter (1989), Van Deemter and Krahmer (1998), and Asher and Lascarides (1998a).

are thus introduced and the situation under discussion. Usually, the discourse contains an antecedent relative to which the anaphor stands in some close, non-trivial relation, for example a relation of functional dependence, or a part-whole relation, or a composite of those.

The following discourse illustrates this. The bridging anaphor 'the driver' links up to the antecedent 'a car crash' by referring to the driver of the car presumably involved in the crash.

(6.18) a. *There was a car crash on the A5. The driver was making a phone-call.*
b. $B \oplus [w_r\, u_3 \mid w_r\colon \textit{car}\ u_3,\ \ w_r\colon \textit{crashed-on-A5}\ u_3]$
c. $[\ \mid w_r\colon \sigma_9\ \ \textit{driver-of}\ \sigma^{\pi}_{10}]$
d. $[u_{20} \mid w_r\colon \textit{car}\ u_{20}] \Rightarrow [\ u_{21} \mid w_r\colon u_{21}\ \ \textit{driver-of}\ u_{20}]$
e. $[\ u_{22} \mid w_r\colon u_{22}\ \ \textit{driver-of}\ u_3]$

For simplicity, we assume that relational nouns such as 'driver' introduce two-place properties such as $\lambda v.\lambda v'[\ |w_r\colon v'\ \textit{driver-of}\ v]$, applied to a self-provided underspecified argument (the thing driven). The latter is represented by the underspecified discourse marker σ^{π}_{10} in the above. In the discourse under consideration (6.18a), the definite description 'the driver' requires that local context (6.18b) entails the presuppositional material (6.18c). By accepting as common knowledge in B that cars have drivers (6.18d), inferring and accepting in B that the car that crashed (u_3) has a driver (6.18d), resolving the incorporated argument of 'driver' to the car in the crash by identifying $\sigma_{10} = u_3$, and resolving 'the driver' to the inferred driver of that car by identifying $\sigma_9 = u_{22}$, the presuppositional condition can be satisfied. Note that in the process a binder for the definite description is accommodated: this is the inferred discourse marker u_{22}.

In the example just discussed, the bridging relation corresponds to the relation expressed by the relational noun, and the link to a specific binder is established by resolving the implicit anaphoric argument of the nominal description. Bridging also occurs with nonrelational nominals, however, which are not normally assumed to contribute implicit anaphoric arguments. Nonetheless, they exhibit the same tendency to link up to elements explicitly mentioned in the immediately preceding discourse, and to 'anchor' to the situation or local setting under discussion. Consider

(6.19) *John got married last April. The priest was bald.*

Zeevat (1992) ex. (42)

'The priest' obviously refers to the individual who performed John's marriage in April. In inferring this link, the hearer may select an implicit background which supports, for instance, that John's marriage in April was a Church marriage, that John is a Catholic, that Catholic marriages are performed by priests. He may then infer that there was a priest involved in perfoming John's marriage, and accept

this information in B. If 'the priest' is bound to this individual, the presuppositional condition is satisfied. Crucially, what is abduced is not just determined on the basis of the descriptive content of the definite description, and general preferences such as the preference for plausible interpretations. The assumption that the priest who is bald is in some way related to John's marriage in April is essential in constraining and guiding the abduction process. Intuitively speaking, there seems to be a restricted domain of linguistic antecedents relative to which the bridging link can be established, and the choice of a candidate antecedent in that domain is the starting point of the inference process.

In general, the inference machinery is quite powerful, and there must be additional constraints at work limiting the abduction process. In line with our earlier observations in section 5.4, we propose that there are additional syntactic, semantic, or informational constraints on definite descriptions which make them resolve *locally*. In particular, definite descriptions, like pronouns, may be taken to require binders that are 'discourse available', given the attachment of the update induced by the sentence they are part of. Definite descriptions put different constraints on discourse availability than pronouns, however. It was tentatively suggested in section 5.4 that while pronouns require their binders to be salient topical entities in the universe of their local context, definite descriptions require some nonsalient form of topicality.

This could be made precise in various ways. It might be assumed that binders of definite descriptions themselves must be topical, as graphically illustrated in (6.20a). The burden of the explanation is then completely on the theory of topicality. In the case of bridging anaphora, *inferred* individual entities or eventualities should count as topical when they are somehow nontrivially and closely linked to topical entities and eventualities that are linguistically introduced in the preceding discourse. The advantage of this is that a generalisation over 'direct binding' and bridging cases is immediately obtained, without having to complicate the semantic analysis of definite descriptions. The difficulty that remains, obviously, is how to delimit the "extended family" of inferrable topical elements that may serve as potential binders given the domain of linguistically introduced topical elements.

(6.20) *definite det:* a.

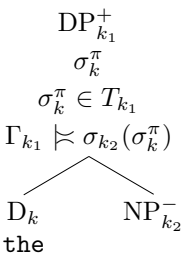

b.

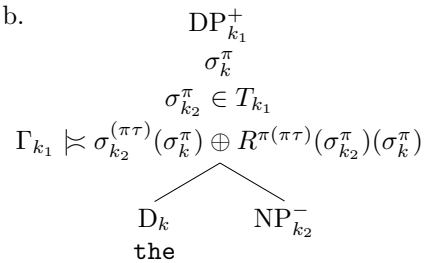

Alternatively, it could be assumed that only individual entities and eventualities that are linguistically introduced can be topical, and that definite descriptions require the linguistically provided entities to which they resolve, either directly

or through bridging, to be topical. To be able to state the latter constraint lexically, the underspecified antecedent marker must be made available. Furthermore, a generalisation over direct resolution and bridging must be provided. The lexical description in (6.20b) indicates a way of doing this. Apart from the descriptive material originating from the nominal phrase $\sigma_{k_2}^{(\pi\tau)}(\sigma_k^{\pi})$, the presuppositional condition contains an underspecified bridging relation R linking σ_k to an underspecified antecedent marker $\sigma_{k_2}^{\pi}$.[22] As before, we are making use of the fact that a single node may come with semantic values of different types. The topicality constraint now says that $\sigma_{k_2}^{\pi}$ must be topical at the attachment site k_1.

The axioms of the grammar require that σ_k and $\sigma_{k_2}^{\pi}$ are bound in the context of interpretation given some value of R. R is a free variable of type $\pi(\pi\tau)$ which is specified in the interpretation process. Consider example (6.19) again. On the assumption that the second sentence in (6.19) is an elaboration of the first, 'John' is a topical element. If we identify the antecedent marker with the discourse marker introduced by 'John' we may find, for instance, $R = \lambda v.w[\ |\ w_r : v\ \mathit{married}\ w\ ]$. By satisfying the presuppositional condition we abduce part of the descriptive content of the definite description; 'the priest' serves as as an abbreviation of 'the priest who married John'.

In principle, any relation that links the antecedent marker and the anaphoric marker given the local context can be abduced. Note that if we find $R = \lambda v.w[\ |\ w_r : v\ \mathbf{is}\ w]$ the definite description is bound directly by the antecedent, thus the treatment generalises over bridging and plain binding. The challenge is to constrain the abduction process. What we want is that solutions of R are non-trivial, not too complex, and intuitively make sense. Part of this may follow from the interaction with the language user's world knowledge, common sense reasoning, and the preference for plausible readings. Another constraining factor is the topicality requirement on the antecedent marker $\sigma_{k_2}^{\pi}$. Since $\sigma_{k_2}^{\pi}$ must be topical, the number of candidates is restricted, and R must 'be about' the element that is chosen, given the hearer's background knowledge. A further restriction may be obtained by requiring that the presuppositional material logically depends on the semantic content of the left argument of the discourse relation that links the triggering sentence to the discourse context. This would ensure that the bridging link that is abduced is *locally justified*. This might be implemented by means of appropriateness conditions associated with the individual discourse operators. We will see more of this in section 6.5.

[22] This is similar to the treatment of bridging relations in Asher and Lascarides (1998b), be it that they do not treat presuppositions in terms of satisfaction requirements on local contexts. Building on the anaphoric approach to presupposition projection, they allow the insertion of presuppositional material in the DRS representing the embedding discourse context. Presuppositional material is linked to the discourse context through underspecified discourse relations. The bridging relation $R^{\pi(\pi\tau)}$ in our treatment corresponds to B in their format.

6.3 Consistency and Informativity Revisited

Presupposition theory is fully integrated with discourse theory in our framework of discourse interpretation. Context-sensitive elements occurring in a discourse may trigger presuppositional conditions. These conditions constrain the discourse meaning, in interaction with independent axioms of the discourse grammar. In this section we zoom in on the interaction with a specific set of axioms and constraints, namely those governing the consistency and informativity of utterances in their context of interpretation. As we shall see, the acceptability conditions on presupposition resolution formulated in Van der Sandt (1988) and (1992) can straightforwardly be implemented by means of consistency and informativity requirements in LDG.[23] We explore the predictions that follow from that treatment a little bit, focusing mainly on the interpretation of conditional sentences. It will be concluded that to obtain a perfect fit between linguistic data and the predictions of the grammar, a further fine-tuning of the theory of informativity is necessary. In particular, the discussion raises the issue whether informativity effects should be handled in the grammar or in the preference system of our framework.

6.3.1 Global Consistency and Global Informativity

As we saw in section 4.4.5 of the preceding chapter, on the basis of independently motivated conversational principles, general conditions on felicitous updating can be formulated and integrated in the discourse theory. Two such conditions were introduced: Global Consistency and Global Informativity. They reflect Stalnaker's (1978) 'first principle' on assertion.[24] The consistency requirement can be stated as a constraint on discourse meanings at object level in our description language. The informativity requirement refers to different stages of the discourse description and must be stated at metalevel in our grammar.

Global Consistency $B \oplus \sigma_r \not\models \bot$

Global Informativity For a discourse contribution updating a discourse description Δ to $\Delta + \delta$, there must be a DRS K such that

$$\Delta^{rcl} \not\models (B \oplus \sigma_r \mid\!\sim K) \text{ and } (\Delta + \delta)^{rcl} \models (B \oplus \sigma_r \mid\!\sim K).$$

Global Consistency says that, for no stage in processing a discourse, the discourse meaning is inconsistent. Global Informativity enforces that every discourse contribution results in a discourse meaning which entails some new information K. These conditions of course interact in a certain way with anaphora resolution and

[23] For an implementation of the acceptability constraints in a computational treatment of the anaphoric approach to presupposition in DRT, see Blackburn and Bos (1999) and Bos (2003).

[24] 'A proposition asserted is always true in some but not all of the possible worlds in the context set.' (Stalnaker 1978), page 325.

presupposition satisfaction, limiting the number of grammatical readings. Global Consistency, for instance, prevents binding relations resulting in readings that are not consistent with information in the general background B. Consider

(6.21) # A^1 *mule is feeding. The man*$_1$ *is content.*

If the language user selects a general background B such that it holds that mules are not men—which is plausible enough because this is general world knowledge—resolving 'the man' to 'a mule' in (6.21) results in an inconsistent discourse meaning. Global Consistency prohibits this, hence, the reading is out.

Furthermore, given some treatment of local accommodation, Global Consistency enforces local accommodation of presuppositions in discourses where global accommodation would result in inconsistency. The following example illustrates this.

(6.22) A: *John made an appointment with Mark's secretary.*
B: *Impossible. Mark doesn't have a secretary.*

Participant A's utterance triggers the presupposition that Mark has a secretary. This is subsequently denied or rejected by participant B. In interpreting participant A's statement an arbitrary hearer[25] may simply accept as general background information that Mark has a secretary, in other words, he may globally accommodate the presupposition.[26] After processing B's reply, however, he finds out that global accommodation is no longer acceptable: the background information that Mark has a secretary contradicts the semantic content of the denial, resulting in an inconsistent discourse meaning. Global Consistency prohibits this.

The situation after B's utterance is that the denial requires a context of interpretation which does not entail that Mark has a secretary, while the presupposition triggered by A's statement must still be satisfied in its local context. The

[25]If participant B is the hearer, things are a bit more complicated. When he evaluates A's utterance relative to his private reservoir of world knowledge $\mathcal{K}$, he finds out that Mark does not have a secretary. So if he accepts the presupposition in the implicit background B, B is not consistent with his world knowledge. We assume that though, preferably, B is consistent with $\mathcal{K}$, it is not required to be so, and there can be stages in the conversation in which it isn't. The hearer may (temporarily) accept information in B for the sake of getting the speaker's meaning, even when he disagrees with it.

[26]Another way of dealing with this type of examples might be to allow the selection of a background such that the presuppositional material is not accepted as a proposotion about the actual world w_r, but about some hypothetical or counterfactual alternative to it, or about an alternative in A's belief space. This obviously requires a specification of an intensional fragment in our discourse semantics. A third option would be to allow greater flexibility in the manner in which presuppositional material is incorporated in the context of interpretation. Beside simply merging it with background information, it could also be added as a condition or hypothesis. This converges with the main idea of Asher and Lascarides (1998b), who propose that presuppositional material is linked to the discourse context through (underspecified) discourse relations.

only way to satisfy *both* requirements while preserving consistency is by allowing some form of local accommodation. If the presuppositional material can be added as local background information to the semantic content of A's statement, the whole of this may be captured in the scope of the negation introduced by the denial.[27] In that case no inconsistency arises and both utterances receive an intuitively appropriate interpretation, relative to a shared global background B. The treatment of local accommodation that will be presented in section 6.4 supports this analysis. Note that in the course of the exchange, the hearer's possible models of the background B change from ones that entails that Mark has a secretary to ones that entail that Mark does not have a secretary. The denial causes a nonmonotonic update of the discourse meaning.

While intuitive judgements about consistency are usually clear and the predictions that follow from Global Consistency are generally correct, informativity is a more complex notion, both in terms of the intuitions that it is supposed to capture and in terms of the formal conditions that should reflect these intuitions. The predictions that follow from Global Informativity are sometimes unexpected, notably where the information growth that may result from global accommodation is concerned. The following three examples illustrate this.

First, observe that when a sentence contains an anaphoric or presuppositional element, Global Informativity blocks readings that result from binding relations which make the sentence uninformative in its context of interpretation. Consider

(6.23) ? A^1 *mule is feeding. The*$_1$ *mule is feeding.*

The question mark indicates our intuitive judgement: the discourse is marked somehow, but not strongly infelicitous. Given resolution of 'the mule' in the second sentence to 'a mule' in the first one, the second sentence is not informative. Global Informativity predicts that this reading is ungrammatical.[28]

Now compare (6.24), where the presupposition triggered by the definite description in the second sentence induces global accommodation and hence information growth. This is a so-called partial matching case, as was discussed in section 6.2. The discourse meaning resulting after the update induced by the second sentence entails that Pedro is a lucky man, which it did not before. Hence, Global Informativity is satisfied, and the second sentence in (6.24) is predicted to be informative.

[27] Compare the lexical description of discourse relation *denial* shown in (6.44) further on. This in fact mimics the treatment of denials proposed in van der Sandt (1991), where the total information content of denied statements ends up in the scope of a negation.

[28] Global Informativity would in principle allow for a reading resulting from globally accommodating the existence of another mule, but that reading should be excluded on independent grounds. We have seen similar examples of this in section 6.2. If we accept a general preference for binding presuppositions to overt antecedents over binding them to accommodated discourse markers, we get the prediction that the reading resulting from global accommodatin is dispreferred. In addition, it may be assumed that topicality requirements on the binder of the definite description, and a general preference for avoiding ambiguity make this reading dispreferred.

(6.24) *Pedro*[1] *owns a mule. The*$_1$ *lucky man owns a mule!*

Finally, observe that if the second sentence in (6.24) is repeated, the resulting third sentence is predicted to be uninformative, and the discourse as a whole ungrammatical:

(6.25) *Pedro*[1] *owns a mule.*
The$_1$ *lucky man owns a mule!*
The$_1$ *lucky man owns a mule!!*

The discourse in (6.25) does not seem to be much less acceptable than the one in (6.24), however. In each case, the second sentence functions as a kind of emotional strengthening of the first, and in the case of (6.25) the third sentence provides additional strengthening. Maybe what one would like to say about these data is that discourse contributions which serve as emotional strengthenings of a statement come with their own kind of informativity. Furthermore, there may be stylistic or literary reasons for repeating information, and this is not necessarily perceived as unfelicitous discourse.

It is open for discussion whether such predictions are desirable or not. As was tentatively concluded in chapter 4, Global Informativity can be assumed to have a limited scope and apply only to discourse contributions which act at the level of plain semantic content, rather than participants' attitudes, and which figure in discourse aimed at information exchange. It could be argued that the continuations in (6.24) and (6.25) serve a different purpose, and are therefore not subject to Global Informativity. But then our explanation of these examples changes: the acceptability of the continuation in (6.24) is not due to the information growth induced by its presupposition, but to its particular rhetorical function in the discourse.

In line with this perspective, it could perhaps be assumed that Global Informativity is a preference, rather than a (hard) constraint in the grammar. As in Asher and Lascarides (2003) it could then be assumed that more specific preferences, or constraints associated with particular discourse relations or linguistic constructs, overrule Global Informativity, when they apply. This would place Global Informativity in the realm of the pragmatics rather than the semantics of our discourse theory. In fact, we propose to take a more radical step, and drop Global Informativity altogether, in favour of local informativity constraints.

As was suggested in chapter 4, the theory of informativity in our grammar can be fine-tuned by means of informativity or appropriateness conditions associated with clause and discourse operators in the lexicon. Whereas Global Informativity is a requirement on the update effect of discourse contributions, local informativity conditions constrain the relation between a discourse unit and its local context. To illustrate, here's a picture of a lexical description of discourse relation *elaboration*, containing a local informativity condition. It is labeled [**inf**].

(6.26)

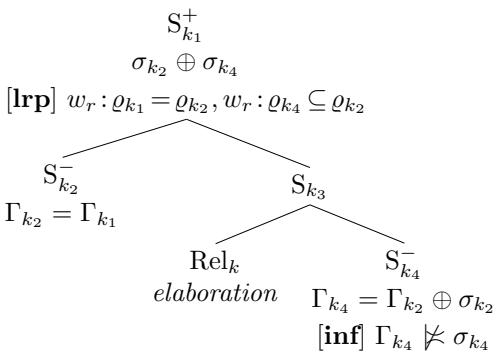

In principle, the individual informativity requirements of particular rhetorical types of discourse contributions can be captured by means of local informativity conditions.[29] Interestingly, in the case of the 'emotional strengthenings' discussed above, it would be the *absence* of local informativity constraints which would characterise these contributions. A treatment with local informativity constraints would predict the continuations in (6.24) and (6.25) to be uninformative, and hence unacceptable, because, unlike Global Informativity, the local constraint would apply to the semantic content of the second sentence, rather than its total update effect.

Since every sentence that serves as the vehicle for a discourse contribution is represented by a discourse unit in the resulting discourse description, the explanatory power of Global Informativity can be maintained by adding local informativity constraints to all discourse operators. Information growth induced by accommodation can be accounted for by either or not setting local informativity constraints on background information. As we saw in section 4.4.5 of chapter 4, however, if we assign local informativity requirements to all S-nodes in a discourse structure the resulting discourse theory easily becomes too strong. The investigations in the next section will confirm this finding, and will lead us to conclude that local informativity conditions can better be treated as 'soft constraints', or constraints that can be violated, in the manner of Optimality Theory.

6.3.2 Van der Sandt's Acceptability Conditions

Conditions on the consistency or informativity of assertions or updates figure in most theories of discourse interpretation. In the anaphoric approach to presupposition projection they are incorporated as *acceptability constraints* on possible

[29] Though only indirectly. The assumption that there is a 1-to-1 correspondence between rhetorical types of discourse contributions and the discourse operators in the lexicon can hardly be maintained in our grammar. In general a composite of discourse operators can be involved in conveying a single discourse contribution. Furthermore, discourse contributions are characterised at meta-level, by their update effect on discourse meanings, whereas discourse operators constrain the compositional semantics of discourse at object level. A characterising feature such as a local informativity constraint of a particular discourse operator is not necessarily inherited by the update as a whole.

solutions, see Van der Sandt (1988) and (1992). Various constraints on binding and accessibility result in a set of 'logically possible interpretations' of a sentence. The acceptibility constraints filter out the contextually acceptable solutions or admissible resolutions among these. Interestingly, Van der Sandt proposes not just global, but also what we would call, in line with Beaver's (2001) discussion of these constraints, *local* conditions on consistency and informativity. They play a crucial role in accounting for cases of presupposition resolution in which accommodation on subordinate levels of DRSs is preferred over global accommodation. Van der Sandt's acceptability constraints are listed in definition 6.1 below.[30]

The definition builds on the following notions and use of terminology. The *incoming DRS* is the DRS representing the discourse processed so far, which is updated when a new sentence in the discourse is interpreted. When a sentence containing one or more anaphoric expressions or presuppositional elements is processed, this results in an unresolved or *preliminary* DRS, containing a so-called A-structure. A-structures are DRSs, representing the semantics of the anaphoric or presuppositional elements in the sentence. An A-structure can be complex, since one anaphoric expression may embed another. Preliminary DRSs are not assigned an interpretation; only when they are merged with the incoming DRS, and when all the anaphoric and presuppositional elements (or A-structures) have been resolved, the resulting DRS is interpretable. We call this a possible *output DRS*.

Anaphora and presupposition *resolution* may take the form of either binding or accommodation. Both operations basically consist in the movement of the 'anaphoric' discourse marker and, when present, the conditions characterising the descriptive content of the anaphor or presupposition, to a target DRS on the accessibility path upward from the position of the anaphoric expression or the A-structure under consideration. In the case of binding, the 'anaphoric' discourse marker is, in addition, identified with an 'antecedent' marker in the universe of the target DRS. To facilitate comparison with our own approach we split out clause (iiia) and (iiib) of the original definition in Van der Sandt (1992), p. 367.

6.1. Definition. **Admissible Resolutions (Van der Sandt 1992)**
Let K_0 be the incoming DRS, K_1 the merge of a DRS with K_0 and K_1' a possible resolution of K_1. The resolution of K_0 to K_1' is subject to the following conditions in order to be admissible:

1. An output DRS K_1' is informative with respect to the incoming DRS K_0, that is, K_0 does not entail K_1'.

[30] A distinction between 'local' and 'global' constraints is made as well in Beaver's reformulation (Beaver 2001) of the acceptability conditions on presupposition resolution and accommodation of Van der Sandt (1992). Beaver's 'Global Informativity' (definition D19, page 107) corresponds to our Global Informativity condition. His 'Local Informativity' corresponds to what we shall call local, or individuated informativity conditions further on, i.e. informativity conditions on local contexts. His 'Consistency' corresponds to what would be consistency conditions on local contexts in our grammar, and covers global consistency as a special case.

2. Resolving K_0 to K_1' maintains consistency.

3. Resolving K_0 to K_1' does not give rise to a structure in which some subordinate DRS K_i is entailed by the DRSs which are superordinate to it.

4. Resolving K_0 to K_1' does not give rise to a structure in which the negation of some subordinate DRS K_i is entailed by the DRSs which are superordinate to it.

The first two conditions in this definition are modeled in LDG by Global Informativity and Global Consistency, respectively. The other two conditions can be stated as follows, on the assumption that 'DRSs that are superordinate to a DRS K_i' correspond to local contexts in LDG.[31]

Local Informativity $\forall k\,[\,\ell(k) = s \rightarrow \Gamma_k \not\models \sigma_k]$

Local Consistency $\forall k\,[\,\ell(k) = s \rightarrow \Gamma_k \not\models [\,|\,\textbf{not}\ \sigma_k]\,]$

This introduces informativity and consistency requirements on the semantic values of all S-clauses or discourse units in a discourse structure. Local informativity and consistency constraints on the arguments of individual discourse connectives and relations can be inferred from these conditions, and integrated in their lexical descriptions. To illustrate, for the connectives 'or', and 'if', this implies the following. Local informativity conditions are labeled by [**inf**], local consistency conditions by [**con**].

(6.27)

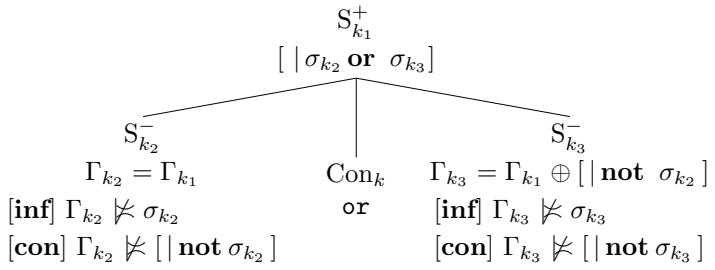

[31] Note that our local contexts have been constructed in the style of Karttunen (1974), so that in some cases the information present in a local context diverges from what would be the corresponding embedding context in standard DRT. For example, with disjunctive connectives or discourse relations, we have that the local context of the right disjunct contains the negation of the semantic value of the left disjunct, whereas in DRT, the left disjunct is simply inaccessible. A full correspondence with DRT may be obtained by slightly changing the construction recipes of local contexts in the descriptions of the relevant discourse operators and connectives.

(6.28)

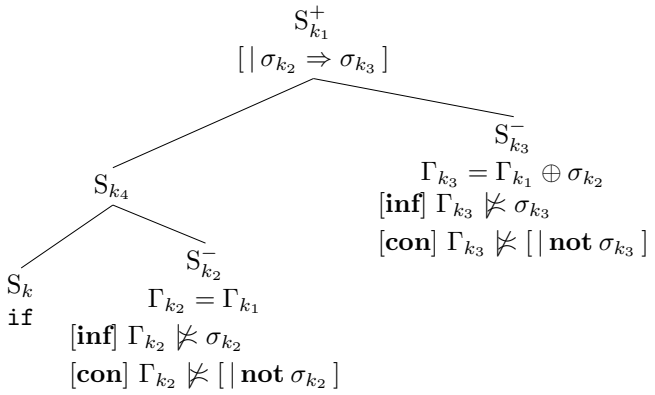

If we want to implement definition 6.1 close to the letter, local consistency and informativity conditions should *not* be assigned to discourse connectives and discourse relations such as 'and' and *elaboration*, whose output semantics is defined just in terms of the merge operator. The conditions 3. and 4. in definition 6.1 apply only when there is a relation of subordination between K_i and the embedding DRS, and in DRT there is no subordination relation between left and right arguments of the merge or conjugation operation. Given the discourse syntactic structure available in LDG, however, our treatment generates local informativity and consistency constraints for the arguments of conjunctive discourse operators as well. We will have, for instance

(6.29)

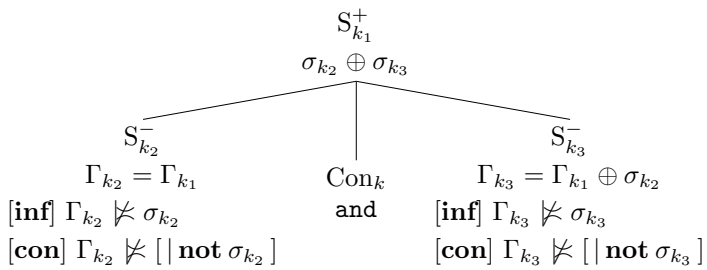

This can be useful in explaining the infelicitousness of sentences in which one of the conjuncts is not informative, given what is presupposed earlier. The following pair of examples illustrates this.[32]

(6.30) a. *Holland has a queen, and the queen of Holland lives in The Hague.*
b. # *The queen of Holland lives in The Hague and Holland has a queen.*

Intuitively, (6.30b) is infelicitous because the right conjunct is not informative in the context of the left conjunct. The local informativity condition on the right argument of the connective 'and' in (6.29) accounts for this. The definite description in the first conjunct triggers the accommodation of the information that there is a

[32]The same point is discussed in Stalnaker (1974), relative to the example 'Harry's wife is great cook, and Harry is married.' See also Beaver (2001) on this topic.

queen of Holland. Relative to the resulting local context, the informativity condition of the right argument of the conjunction is not satisfied. Hence, the discourse is predicted to be infelicitous. In contrast, in (6.30a) the informativity condition of the right argument of the conjunction *is* satisfied. A similar argument can be made in favour of the local informativity condition on the left argument. As noted before, however, if local informativity conditions are accepted in conjunctive constructions the resulting discourse theory easily becomes too strong, due to the interaction with other appropriateness constraints.

The implementation of Van der Sandt's acceptability conditions in LDG opens up the possibility to investigate these requirements as integrated in a more encompassing discourse theory. For a start, we may ask what intuitive notion of 'informativity' is captured by local informativity conditions. Global Informativity expresses a felicity condition on assertions, that is, discourse contributions of a certain rethorical type, and it is reasonably clear what sort of intuitions underly that requirement. But what do local informativity conditions stand for? As Van der Sandt puts it, in incremental terms one could say that no provisional update may lead to inconsistency and each provisional update should at least provide some new information. Correspondingly, the conditions in definition 6.1 mark for instance the pieces of discourse in (6.31) as unacceptable.

(6.31) a. ? *John has a dog. If he has a dog, he has a cat.*
b. ? *John has a dog. If he has a cat, he has no dog.*
c. ? *John has no dog. Either he has a dog or he has a cat.*

Van der Sandt (1992), ex.(63a-c), p.368

The informativity and consistency conditions attached to the connectives 'if' and 'or' in our grammar predict these texts to be ungrammatical: (6.31a) violates local informativity on the antecedent of the conditional, (6.31b) violates local consistency on the consequent of the conditional, and (6.31c) violates local informativity on the left argument of the disjunction. According to our intuitions, however, the texts in (6.31) are just 'marked' or slightly infelicitous, when uttered out-of-the-blue. Thus, the local informativity and consistency conditions are more constraining than our intuitive judgement of the data warrants.

As Van der Sandt observes, a local violation of consistency or informativeness need not give rise to uninformativeness of the whole utterance processed. It often signals that the information carried by the utterance is conveyed in an unnecessarily redundant and complex way. If, for example, ϕ has already been established in the context of interpretation, a statement of the form 'if ϕ then ψ' is clearly informative, since it tells us that ψ is the case. However, the same information could have been conveyed in a shorter and thus more efficient way simply by stating ψ. Efficiency and (global) informativeness thus are distinct notions, Van der Sandt concludes.

If this is on the right track, it is questionable whether violations of efficiency deserve the heavy penalty they are getting now in our grammar, namely the prediction that the resulting discourse is ungrammatical. A less stringent treatment, which would predict that these cases are *dispreferred* rather than ungrammatical, might be more appropriate. But this would mean that local informativity constraints should not be treated on a par with the other (syntactic and semantic) constraints in the grammar. They could perhaps be made to follow as 'soft constraints' from a general preference for more efficient formulations over less efficient ones. **Local Informativity** and **Local Consistency** could be seen as the embodiment of that preference.

Setting this matter aside for the moment, let us take a look at some more examples. The predictions that follow from local informativity conditions in the grammar generally do better on presuppositional information than on semantic contents, and in fact constitute the main reason for incorporating them in the discourse theory. As Van der Sandt argues, there are classes of examples in which the only admissible readings of a sentence are non-presupposing ones,[33] and this can be explained in terms of the local consistency and informativity constraints. These examples are characterised by the fact that what is presupposed by one argument of a connective is either entailed or contradicted by the other argument in the context of interpretation. Consider the following case, which is a variant of example (69) in Van der Sandt (1992).

(6.32) *If John*^{1} *is married, he brings his*$_1$ *wife*^{2}*.*

Given resolution of the possessive 'his' to 'John', the definite description triggers the presuppositional condition that John has a wife. Intuitively, a hearer does not conclude from (6.32) that John is married, in other words, the conditional as a whole is not presupposing. Given the local informativity condition associated with the antecedent of the conditional, this is precisely what the grammar predicts. The condition prevents global accommodation: if it were globally accommodated that John has a wife, the informativity requirement on 'John is married' could not be satisfied. The only reading allowed by the grammar is a non-presupposing one, which can be paraphrased as "If John is married, he has a wife and he brings her" or as "If John is married and has a wife, he brings her". In the absence of the local informativity constraint, the grammar would predict there is a possible presupposing reading as well. In a similar way, nonpresupposing readings of disjunctive statements can be accounted for in terms of local consistency constraints attached to the disjunctive operator.

Notwithstanding these positive effects of incorporating local informativity and consistency constraints, the overall strength of the grammar continues to be a problem, if local informativity and consistency conditions are attached freely to

[33] A *non-presupposing* reading of a sentence results when the presuppositional material in the sentence is not globally bound or accommodated, in the Van der Sandtian sense of these terms.

all available arguments of the relevant discourse and sentence operators, and figure as hard constraints rather than preferences. Consider

(6.33) *If John*^{1} *brings his*$_1$ *wife*^{2}*,* *he*$_1$ *is married.*

The conditional is perfectly acceptable in a nonpresupposing reading which can be paraphrased as "If John brings a person who is his wife, then he is married". This reading is predicted to be unacceptable, however, because the local informativity condition on the consequent of the conditional cannot be satisfied: John's having a wife implies that he is married. The fact that the reading exists is not surprising, since conditionals generally serve to express that if (in the given context) the antecedent holds, the consequent follows. We conclude that at least for certain uses of the conditional, the local informativity requirement on the *consequent* of a conditional should be dropped, or allowed to be overruled.[34] In fact, the local informativity constraint on the *antecedent* of the conditional is also unwanted in certain cases. The following discourse illustrates this.

(6.34) *George will be present at the party,*
and if George will be there, I'm not coming.

Intuitively, the conditional in (6.34) is felicitous in the context in which it is uttered. The speaker conveys that she is not coming to the party, and is making explicit the reasoning behind this decision. The discourse is predicted to be unacceptable, however, because the local informativity condition on the antecedent of the conditional is violated. We have seen this type of example before in (6.31a). It was suggested then that violations of local informativity conditions predict inefficiency rather than uninformativeness of the statement under consideration. The example (6.34) does not seem particularly inefficient to us, however, or infelicitous in any other way. It simply instantiates one of the possible uses of conditionals in discourse. In section 4.4.5 of chapter 4, a more extreme instance of this use was mentioned, namely the use of a conditional statement in the context of a proof by modus ponens. There, the antecedent is typically not informative; it might even be claimed that it is required to be *un*informative.

6.3.3 Conclusion - Ways to Go

Starting out from the idea that a theory of informativity can be incorporated in the grammar of our framework of discourse interpretation, we specified two

[34] We might even consider replacing the informativity and consistency constraints by a requirement that expresses the consequential relation inherent in conditionals, [**csq**] $\Gamma_{k_3} \mid\!\sim \sigma_{k_3}$. As was argued in section 4.4.5 of chapter 4, however, the condition [**csq**] easily gets into conflict with informativity constraints. In fact, if we adopt [**csq**] the only reading of (6.33) allowed by the grammar is one in which the local [**csq**] requirement induces global accommodation of the conditional statement, and hence information growth. So perhaps it is better to refrain from the [**csq**] condition.

global constraints on discourse meanings, and proposed that the resulting theory of informativity can be refined by adding local informativity and consistency conditions to elements in the lexicon. In fact, we proposed to replace the metalevel Global Informativity condition altogether by (object level) local constraints.

As we saw, the acceptability conditions on sentence interpretation formulated in Van der Sandt (1992) can be implemented by means of informativity and consistency conditions in our grammar. Testing the predictions of the theory, we got a rather mixed result: in some cases the predictions fit our intuitive judgement of the data, in others they exclude discourses which intuitively speaking are acceptable. This suggests that our treatment is in some respects too restrictive or not flexible enough. We can think of some different ways of obtaining a better fit between the linguistic observations and the formal treatment. We mention them here without making a definite choice, further research will have to show what the most succesful approach will be.

Firstly, we can escape the unwanted predictions by assuming that local informativity constraints, when they obtain, are not introduced in the lexicon, but are 'soft constraints' which follow from general preferences on, among others, the informativity of assertions, and efficiency of formulation. This type of account clearly calls for a further specification of the preference system. It might for instance be assumed that a preference for efficient formulation is counterpoised by a preference for easy traceability or 'perspicuity' by the hearer. One would still have to explain why in some exchanges the preference for efficiency wins over the preference for traceability, while in cases such as (6.34) it does not. The implication of this solution is that the theory of informativity is a matter of pragmatics rather than semantics. The predictions of this type of account will be more flexible, but also much less clearcut, as long as the preference system itself is open ended. As for a practical implementation, perhaps a treatment with violable constraints along the lines of Optimality Theory could be provided.[35]

An alternative solution would be to assume that different *uses* of constructs such as conditionals come with different local informativity constraints, and to underspecify lexical descriptions of conditionals for these different uses. A disambiguating feature for 'text-genre' or 'use' such as logical proof, or plain information exchange, might be added to the grammar, and discourse operators made sensitive to this feature. This solution only has explanatory power if, in general, discourse contexts are sufficiently specific or informative to disambiguate the different uses or genres. An important factor in disambiguating between different informational variants of connectives may well be the topic structure of discourse. The selection of one specification or another of the connective could be linked to what is the local discourse topic, or topmost issue under discussion, given the attachment of the connective or discourse relation.

[35]See Hendriks and de Hoop (2001) and Blutner and Zeevat (2004), for an entrance to the literature on semantics and pragmatics in OT.

Finally, we could drop the local informativity constraints as they were formulated in the above, and try to provide alternative conditions that are better tuned to account for the data under consideration. This is the line of reasoning taken in Asher and Lascarides (1998b) and (2003), where specific requirements associated with the use of conditional constructions and discourse relations, in interaction with the update operation as such, and their principle of Maximising Discourse Coherence account for the data.[36] The specific requirements associated with conditionals essentially capture the relation of logical dependence between antecedent and consequent. If the local incorporation of presuppositional material contributes to the naturalness of the relation of dependence rather than not, it is predicted to be preferred over global insertion. A treatment like this could in principle be modeled in LDG by associating appropriateness conditions that capture the dependency with conditional connectives or discourse operators.

6.4 Generalised Accommodation

As it stands, our discourse grammar handles only global accommodation or specification of the implicit background of the conversation B. The hearer may accept general information in the implicit background, for the purpose of obtaining a suitable interpretation of a newly processed utterance in the ongoing discourse. The background must support what is presupposed, and the discourse meaning as a whole must be plausible given the hearer's world knowledge, common sense reasoning and preferences. As in the anaphoric approach to presupposition projection, information which is globally accommodated ends up in the main DRS of the discourse meaning, and is entailed by it. Unlike that approach, however, what is globally accommodated need not be identical to the presuppositional material of a given trigger, it can be any piece of world knowledge or general knowledge.

Although it can perhaps be argued that global accommodation is all that is needed in presupposition theory, see for instance Van Rooij (2005), most theories employ forms of local and intermediate accommodation, in order to account for presuppositional sentences whose presuppositions do not project to the level of the main context. *Local* and *intermediate* accommodation may be defined as accommodation relative to some local context other than the initial or main context of interpretation such that the accommodated information is not entailed by the resulting discourse meaning. In the anaphoric approach to presupposition what is accommodated locally or intermediately ends up in a subordinated DRS, rather than the main one. *Local* accommodation is accommodation which results in the expansion of the *smallest* DRS that contains the condition(s) in which the semantic value of the anaphor or presupposition trigger figures. Accommodation

[36] Asher and Lascarides argue, on the basis of a different set of data then discussed here, that the informativity and consistency requirements are somewhat too weak to account for projection facts.

resulting in the expansion of a DRS strictly subordinating the smallest one, other than the main one, is called *intermediate* accommodation.

In this section we put forward an extension of the grammar which allows for local and intermediate accommodation as well as global accommodation. Intuitively speaking, the idea that background information may occasionally help to strengthen the semantic content of a discourse rather than constrain the initial context of interpretation is not implausible. Presuppositions may constrain local contexts other than B, and say something about situations or states that are in a way subordinate to the actual state of affairs the discourse as a whole is assumed to be 'about'. From this point of view, it seems natural to attempt a generalisation of our treatment to accommodation within σ_r. We shall call this *content modification*. From a practical point of view, we simply want to facilitate the integration in LDG of approaches to presupposition theory that employ local, or local and intermediate accommodation.

After an informal discussion of some basic cases, in which intuitively speaking, presuppositions do not project to top level (section 6.4.1), we will turn to an implementation of content modification in LDG (section 6.4.2). The basic idea will be that accommodation as such need not be restricted to the implicit general background B alone, but that all discourse units in a discourse tree structure come with their own local background, and that these local backgrounds become specified as a consequence of the interpretation process. To avoid unwanted weakening of discourse meanings as a result of partial underspecification of local backgrounds, they shall be required to be fully specified and constructed of linguistically generated material only. Thus, content modification is resource bound. As such, it closely corresponds to the constructive notion of accommodation of Van der Sandt (1992) c.s. In fact, a full correspondence with the anaphoric account of presupposition projection can be obtained by making B resource bound as well. Since we are interested in combining the Karttunen (1974)-style treatment of presupposition with the anaphoric approach we will not do that here, however.

In the resulting treatment, a given presupposition can be satisfied via accommodation of the presuppositional material in any one of the local backgrounds on the accessibility path from the trigger upwards to the main DRS of the discourse meaning, including the local background of the root node. In principle, each of these options constitutes a possible reading of the discourse containing the trigger. At the same time, further specification of the general background B may take place. Global accommodation and content modification come out as qualitatively different forms of accommodation, which can co-occur. Overgeneration is avoided through interaction with the interpreting agent's preference system and common sense reasoning. In section 6.4.3 various preferences are discussed. This includes the central preferences in the anaphoric approach to presupposition, as well as the preference for plausible readings commonly adopted in the satisfaction approach. In addition, we will propose a preference for minimizing local backgrounds, on

the basis of a general economy principle.

6.4.1 Some Examples of Nonpresupposing Discourse

We have already seen some examples in which, intuitively speaking, the most preferred reading of a possibly complex sentence is a nonpresupposing one. That is, a reading in which the presupposition under consideration is not bound and satisfied in the context of interpretation of the whole sentence, or made to be so through global accommodation.[37] Nonpresupposing readings are the ones that call for some form of local or intermediate accommodation, if accommodation is required. Consider the following basic case.[38]

(6.35) *If John*^{1} *is married, he brings his*$_1$ *wife*^{2}*.*

Intuitively speaking, when uttered in an out-of-the-blue context, the most preferred reading of the conditional in (6.35) is a nonpresupposing one; we do not conclude from (6.35) that John has a wife. The reason for this is that, given world knowledge and common sense reasoning, John's having a wife *depends* on John being married or not. On the assumption that the latter issue is undecided as yet, it cannot be concluded that John has a wife. The most preferred readings are ones in which the presupposition is bound and satisfied in the scope of the conditional. They can be paraphrased as "If John is married, he has a wife and he brings her" and "If John is married and has a wife, he brings her."

According to our intuitions, a reading in which we accept globally that John is married and has a wife, is not impossible but strongly dispreferred in an out-of-the blue context.[39] In the spirit of the satisfaction approach to presupposition, it may be assumed that the preference for the nonpresupposing reading results from the interpreting agent's reasoning about world knowledge and a preference for selecting most plausible discourse meanings. In line with the anaphoric approach, the preference can be explained in terms of a local informativity condition associated with the antecedent of the conditional.

[37] We are using terminology in the sense of our own treatment here, in which presuppositional conditions must be satisfied in their local context *given* binding of the discourse markers free in them. Anaphoric presupposition triggers are always bound, either by a linguistic antecedent or by accommodated material. In Van der Sandt's approach presuppositions are *either* bound *or* accommodated.

[38] Note that in the satisfaction approach, the presupposition triggered by 'his wife' in (6.35) is satisfied directly in the local context established by the antecedent of the conditional, given the resolution of 'his' to 'John'. No local or intermediate accommodation is necessary. The reason is that the presuppositional material is existentially bound rather than anaphoric. In our combined approach, however, the presuppositional material is anaphoric, that is, the main discourse marker is free in it. Hence, a binder must be accommodated in order to satisfy the presupposition.

[39] A context of utterance in which this reading comes out as the preferred one is one in which the conditional is part of a modus ponens argument. This use of conditionals was touched upon in section 6.3.

If there is no logical dependence between the antecedent of the conditional and the presupposition, the presuppositional material may escape the scope of the conditional. This is illustrated by the following example. In this case, the most preferred reading of the conditional in an out-of-the-blue context is a presupposing one.

(6.36) *If John^1 is coming to the conference, he brings his$_1$wife^2.*

We prefer a reading of (6.36) from which it can be concluded that John has a wife. Given world knowledge and common sense reasoning, John's having a wife does not depend on his coming to the conference. Since many men are married, it may simply be accepted globally that John has a wife. According to our intuitions, there are two nonpresupposing possible readings, which are both dispreferred in an out-of-the-blue context. One of these can be paraphrased as "If John is coming to the conference he has a wife and he brings her." This is an implausible reading because it implies that John's having a wife somehow depends on his coming to the conference.

The other nonpresupposing reading can be paraphrased as "If John is coming to the conference and he has a wife, he brings her." It cannot very well be argued that this is an implausible reading. Still, it is clearly less preferred than the presupposing reading of (6.36). In line with the anaphoric approach, the preference for the presupposing reading may be accounted for in terms of a general preference for global accommodation over local or intermediate accommodation. Alternatively it could perhaps be explained in terms of specific appropriateness conditions associated with conditionals. It might be assumed that presuppositional material is only accommodated in the scope of a conditional if it nontrivially depends on the antecedent, given the context of interpretation.

The idea that a theory of presupposition requires a treatment of *intermediate* accommodation is often motivated on the basis of examples of the following kind.[40]

(6.38) *Either Wilma isn't married,*
or she believes that her husband is deceiving her.

Geurts and van der Sandt (1998), ex. (32b)

(6.39) *If a rich man crosses the city-centre,*
he enjoys doing it in his fourwheel-drive.

[40] Additional evidence is provided in van der Sandt (2007), who discusses Kripke's example

(6.37) *The number of planets might have been necessarily even.*

Kripke (1979) ex. (41)

where the interaction with background knowledge blocks both the 'global' and the 'local reading'. Only the reading resulting from intermediate accommodation—in the scope of the might-operator but outside the scope of the necessity-operator—is acceptable.

Consider (6.38). The definite description 'her husband' introduces a presuppositional condition which must be satisfied in the local context established by the disjunction and the belief operator in the second disjunct. Suppose we resolve 'she' and the two occurrences of 'her' to Wilma. Intuitively speaking, (6.38) is not presupposing, that is, we do not conclude from the whole of (6.38) that Wilma has a husband. The sentence has two nonpresupposing possible readings, which can be paraphrased as in (6.40a) and (6.40b).

(6.40) a. *Either Wilma is not married, or (she is married and) she has a husband and she believes that he is deceiving her.*

b. *Either Wilma is not married, or she believes that (she is married and) she has a husband and he is deceiving her.*

We assume the material in parentheses is inferred in the interpretation process, but not necessarily accommodated. In (6.40b) the presupposition is satisfied by accommodating it at the deepest embedded available location, namely in the scope of the belief operator in the right disjunct. This is local accommodation. In (6.40a) the presupposition is satisfied by accommodating it one level less deeply embedded, namely in the scope of the disjunction but outside the scope of the belief operator. This is intermediate accommodation. To our taste, the 'intermediate reading' (6.40a) is the most preferred one of the two.

Setting aside the question what the exact semantics of the belief operator must be, the preference for a nonpresupposing reading over a presupposing one can readily be understood. If—without taking a counterfactual or revisional point of view—it is considered possible that Wilma is not married, it cannot be accepted that she has a husband, because that would imply she is married. Global accommodation would result in an inconsistent local context, and hence an implausible reading. In line with the anaphoric account of presupposition, local consistency requirements can be associated with the individual disjuncts in order to prohibit this. The preference for 'intermediate' (6.40a) over 'local' (6.40b) may perhaps be explained in terms of the overall preference for more plausible readings over less plausible ones, or, in line with Van der Sandt (1992) and Geurts (1999), in terms of a preference for accommodation as closely as possible to the main DRS.

In the other example, here repeated in (6.41a), again, the most preferred reading is often claimed to be one resulting from intermediate accommodation of the presupposition triggered by 'his fourwheel-drive'. We consider only readings where the possessive pronoun 'his' resolves to 'a rich man' in the antecedent of the conditional. Given that restriction, properness requirements exclude a presupposing reading resulting from global accommodation. We consider two nonpresupposing readings, one resulting from intermediate accommodation and paraphrased in (6.41b), the other resulting from local accommodation and paraphrased in (6.41c).

(6.41) a. *If a rich man crosses the city-centre, he enjoys doing it in his fourwheel-drive.*

b. *If a rich man owning a fourwheel-drive crosses the city-centre, he enjoys crossing it in that car.*

c. *If a rich man crosses the city-centre, he owns a fourwheel-drive and enjoys crossing it in that car.*

(6.41b) is a statement about (the subset of) rich men who own a fourwheel-drive. The reading can be considered plausible if the individuals under discussion can be taken to be 'rich men who own fourwheel-drives'. (6.41c) is a statement about rich men, and implies that all rich men who cross the city-centre own fourwheel-drives. This reading can be considered plausibe if it is accepted as general background information that rich men normally own fourwheel-drives and if the individuals under discussion can be taken to be default cases of rich men. Considering (6.41a) in an out-of-the-blue context, reading (6.41b) seems to be preferred over reading (6.41c). The anaphoric approach explains the preference for 'intermediate' (6.41b) over 'local' (6.41c) in terms of a preference for accommodation as closely to the main DRS as possible, as before. An explanantion in terms of a preference for most plausible interpretations alone does not seem sufficient, though it could possibly be argued that what needs to be accommodated to obtain a plausible reading of 'local' (6.41c) is overall more costly than is the case with 'intermediate' (6.41b).

Intuitions about 'intermediate' versus 'local' readings are, generally speaking, subtle, and often disputed among authors. There is an ongoing debate about these and similar data in relation to presupposition accommodation, domain restriction in conditionals and quantificational constructions, and the topical structure of discourse (Beaver 1994; Asher and Lascarides 1998b; Geurts and van der Sandt 1998; Asher and Lascarides 2003). We will not decide in this matter. However, since we want our discourse theory to cover nonpresupposing readings like the ones discussed above, we will make available the technical means to accommodate locally and intermediately.

6.4.2 Content Modification— Local, Intermediate, and in the Main DRS

We start out from the idea that accommodation need not be restricted to the general background B of a discourse, but may apply to its semantic content σ_r as well. We propose that accommodation may occur at any S-node k in the discourse description and strengthen its semantic content with background information. In particular, each S-node comes with an underspecified *local background*—beside its local context—that may absorb presuppositional material. Local backgrounds shall be merged with the semantic values of the nodes. The semantic content of a discourse will be composed of the resulting values and local contexts will be constructed in terms of them.

We must take care now that our notion of discourse meaning does not become to weak. Discourse meaning, in our underspecification formalism, is essentially what can be deduced from $B \oplus \sigma_r$ given the discourse description, the grammar, and the hearer's world knowledge and preferences. Having underspecified local backgrounds as components in the compositional semantics of discourse leads to a weakening of the discourse meaning whenever an underspecified local background sits in the scope of a logical operator other than plain conjunction. The interpreting agent cannot draw the conclusions which intuitively speaking, he should be able to draw. To remedy this, we require that local backgrounds are resource bound, or more precisely, designated to be composed from linguistically generated material only. Though local backgrounds are underspecified, the language user knows or can deduce what they are maximally composed of. The possible specifications of a given local background can be summed up.

This can be implemented in the grammar in the following way. Every S-node k in a discourse description, including the root node r, comes with a *background set* $\mathbf{b}(k)$ of type νt, a possibly empty set of nodes. The nodes function as pointers to linguistically generated material. In principle, any lexical element or implicit anchor with accommodation potential can provide this material, but we focus on anaphors and presupposition triggers here. These come with a value Π_k of type τ which constitutes the *background material* generated by the trigger. In the case of anaphoric triggers, it is composed of the presuppositional material plus a binder of the anaphoric marker of the trigger. To illustrate, (6.42) shows updated versions of the lexical descriptions of the definite determiner 'the' and the pronoun 'he'. For elements in the lexicon with no accommodation potential, the Π-value must be set to [|].

(6.42) *definite det:*

$$\begin{array}{c} \mathrm{DP}^{+}_{k_1} \\ \sigma^{\pi}_{k} \\ \sigma_k \in T_{k_1} \\ \Gamma_{k_1} \mid\!\sim \sigma_{k_2}(\sigma^{\pi}_{k}) \\ \Pi_k = [\sigma^{\pi}_{k} \mid\,] \oplus \sigma_{k_2}(\sigma^{\pi}_{k}) \\ \diagup \quad \diagdown \\ \mathrm{D}_k \qquad \mathrm{NP}^{-}_{k_2} \\ \texttt{the} \end{array}$$

pronoun:

$$\begin{array}{c} \mathrm{DP}^{+}_{k_1} \\ \sigma^{\pi}_{k} \\ \sigma^{\pi}_{k} \in T^{mostsalient}_{k_1} \\ \Pi_k = [\sigma^{\pi}_{k} \mid\,] \\ | \\ \mathrm{D}_k \\ \texttt{he} \end{array}$$

When an underspecified binder σ_k is accommodated, a discourse marker that is *new* in the domain of discourse must be chosen.

For every node k, a *local background* or *background value* β_k of type τ is computed from its background set $\mathbf{b}(k)$ by taking the (generalised) merge of the background material at the nodes in $\mathbf{b}(k)$:[41]

[41] In this, and the following axioms, $\forall k, k'[\, k' \in \mathbf{b}(k)]$ stands for $\forall k, k'[\, \mathbf{b}(k)(k')]$, and $\forall k[\, \beta_k = \bigoplus\{\Pi_{k'} \mid k' \in \mathbf{b}(k)\}\,]$ abbreviates the following statement:

$\mathcal{A}23\ \forall k[\,\beta_k = \bigoplus\{\Pi_{k'} |\, k' \in \mathbf{b}(k)\}\,]$

In order to account for intermediate accommodation in conditionals and quantifier structures, it may be assumed that S-nodes come, in addition, with a subsidiary background set $\bar{\mathbf{b}}(k)$ of type νt, likewise containing pointers to background material. Axiom $\mathcal{A}23$ must then be strengthened to $\mathcal{A}23'$.

$\mathcal{A}23'\ \forall k[\,\beta_k = \bigoplus\{\Pi_{k'} |\, k' \in \mathbf{b}(k)\} \wedge \bar{\beta}_k = \bigoplus\{\Pi_{k'} |\, k' \in \bar{\mathbf{b}}(k)\}\,]$

In any given discourse tree structure, local backgrounds must be composed of presuppositional material generated by the stretch of discourse they are the local background of. The following axiom takes care of this.

$\mathcal{A}24\ \forall k, k'[(k' \in \mathbf{b}(k)\ \vee\ k' \in \bar{\mathbf{b}}(k)) \rightarrow (term(k)\ \wedge\ k \lhd^* k')]$

Axiom $\mathcal{A}24$ says that background sets $\mathbf{b}(k)$, or $\bar{\mathbf{b}}(k)$, contain only terminal nodes k' dominated by k. The effect of this is that the background value of an arbitrary node or discourse unit k will be constructed only from background material introduced by elements in the lexical yield of k. All of the material in it can be traced back to lexical elements or implicit anchors occurring in the relevant stretch of discourse.

Arguably, background material, if it is accommodated, can only be consumed once. To avoid multiple accommodation, axiom $\mathcal{A}25$ may be adopted.[42]

$\mathcal{A}25\ \forall k, k', k''[(k' \in \mathbf{b}(k)\ \wedge\ k' \in \mathbf{b}(k'')) \rightarrow\ k'' = k]$ and
$\forall k, k', k''[(k' \in \bar{\mathbf{b}}(k) \leftrightarrow k' \notin \mathbf{b}(k)) \wedge ((k' \in \bar{\mathbf{b}}(k)\ \wedge\ k' \in \bar{\mathbf{b}}(k'')) \rightarrow\ k'' = k)]$

Another option would be to adopt a preference for smaller, or smallest possible local backgrounds. This will be discussed further on. The preference would result in the prediction that cases of multiple accommodation are dispreferred in general, though not excluded under all circumstances. If that prediction sufficiently fits the data, $\mathcal{A}25$ could be dropped in favour of the preference.

Background values are to be incorporated in the semantic content of discourse, and in the local contexts of discourse constituents. We will do this by defining the output semantic value of sentence and discourse operators in terms of the 'information content' of their arguments, where information contents are composed of background values and semantic values. For each S-node k, its *information content* i_k is the merge of its background value and its semantic value:

$\forall k[\forall v(\mathsf{fst}(\beta(k))(v) \leftrightarrow \exists k'(\mathbf{b}(k)(k') \wedge \mathsf{fst}(\Pi(k'))(v))) \wedge$
$\forall c(\mathsf{snd}(\beta(k))(c) \leftrightarrow \exists k'(\mathbf{b}(k)(k') \wedge \mathsf{snd}(\Pi(k'))(c)))]$

[42] But compare Zeevat (1992), who allows for accommodation of a single presupposition at all levels of embedding downward from the level of accommodation maximally allowed by the other constraints on the interpretation process, and who employs multiple accommodation to account for presupposition projection in belief contexts.

$\mathcal{A}26\ \forall k[\ell(k) = s \rightarrow i_k = \beta_k \oplus \sigma_k]$

Discourse meanings may now be described as the update of the general background B with the information content of the discourse: $B \oplus i_r$.

As we said, the output semantic values of sentence and discourse operators are composed of the information contents of their arguments, rather than their bare semantic values. In the case of conditional connectives or discourse relations, and quantificational operators other than the indefinite, they also employ subsidiary background values.[43] Likewise, local contexts incorporate background values: for any given S-node k, its local context is constructed top down from the local context of its mother node, the information content of left sister nodes when present, and its own local background. All of this is laid down in the lexicon of the grammar, as illustrated by the lexical descriptions in (6.43), (6.44) and (6.45). In addition, the following axiom, which is a revised version of $\mathcal{A}18$, sets the local context of the root node r to $B \oplus \beta_r$ and the local contexts of non S-nodes to the local context of their mother node.

$\mathcal{A}18'\quad \Gamma(r) = B \oplus \beta_r\ \wedge\ \forall k_1 k_2\,[[k_1 \lhd k_2 \wedge \ell(k_2) \neq s] \rightarrow \Gamma(k_2) = \Gamma(k_1)]$

(6.43)

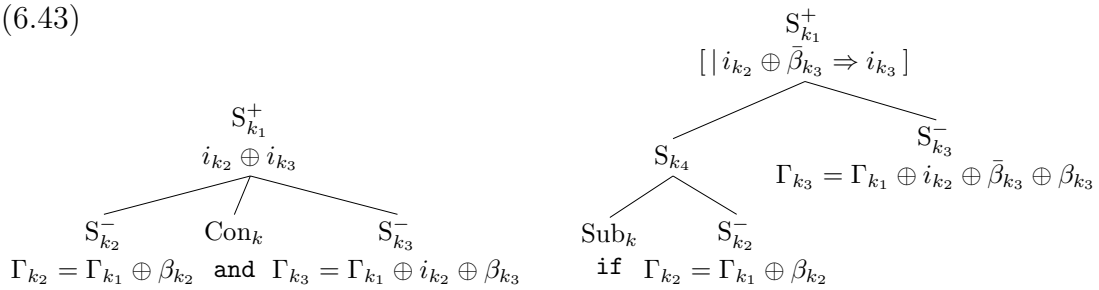

[43] These operators are often claimed to allow intermediate accommodation in their antecedent or restrictor. The relation between the antecedent or restrictor and the presupposition trigger cannot be described in our grammar purely in terms of the syntactic relation of dominance. A description in terms of c-command would be feasible but would enforce a different syntactic form on quantifiers than the one adopted in LDG. The subordination relation in DRT could then be seen as a relation between nodes in the discouse tree structure and defined in terms of c-command and dominance. However, we decided not to change the syntax of quantifiers in our grammar and employ subsidiary background values instead. For those who believe that 'intermediate readings' of conditionals and quantifiers can, technically speaking, be obtained through other means than intermediate accommodation of presuppositional material, the grammar formalism can easily be adapted: the subsidiary background values and the axioms in the grammar concerning them can simply be excluded or pruned.

(6.44)

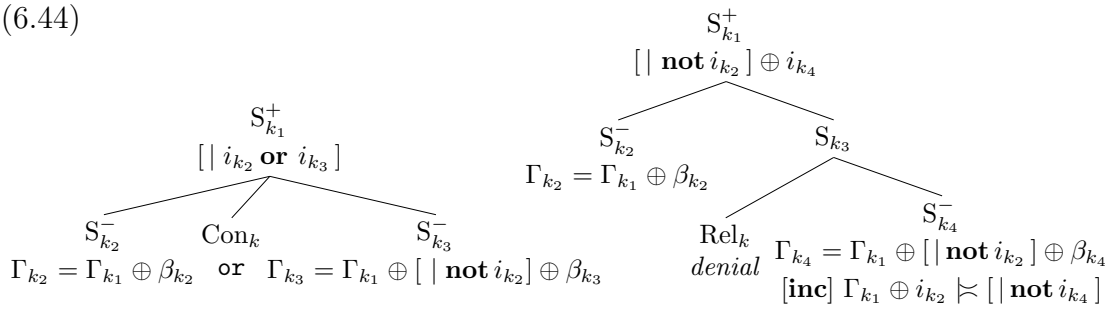

(6.45)

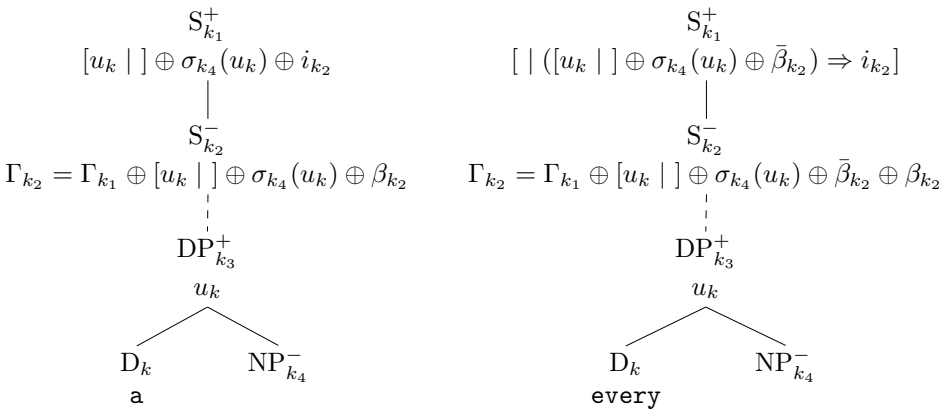

With the above definitions and revised lexical descriptions put in place, the information content of a discourse i_r can be modified or locally strengthened in the interpretation process. As before, presuppositional elements contain a condition which says that the presuppositional material must be entailed by the local context of the element. That requirement may be satisfied directly, by binding the presupposition trigger to a discourse marker which was already introduced by an antecedent present in the discourse. It may also be satisfied by accommodating the presuppositional material in one of the background values in the local context of the trigger. In that case, the trigger is bound by the accommodated material. All of this may co-occur with global accommodation, that is, the specification of world knowledge in the implicit background to the conversation B.[44]

[44]In a brief paper, Hunter and Asher (2005) propose a treatment of presupposition, which, like ours, is hybrid in the sense that it combines the possibility to satisfy a presupposition with the option to accommodate or insert presuppositional material. As they say 'presupposition is resolvable in a sequence of contexts just in case some subsequence entails the presupposition or it is accommodated at some element in the sequence. Resolving the presupposition means choosing some witness for the existential quantifier that binds the presuppositional material.' Satisfaction-plus-choosing a witness results in what we have called 'binding' readings in the anaphoric approach. Insertion or accommodation of the presuppositional material corresponds to 'accommodation' in the anaphoric approach. Their set-up seems to differ in perspective from ours in the sense that the starting point of our treatment is that presuppositional conditions are always there and must be satisfied. This may or may not involve accommodation. In our treatment, accommodation and satisfaction are not viewed as alternative resolution procedures.

While global accommodation implements the concept of accommodation of the satisfaction approach to presupposition, content modification corresponds to the notion of accommodation in the anaphoric approach. Our treatment of content modification compares to the anaphoric approach in the sense that the readings resulting in that approach from moving presuppositional or anaphoric material along the accessibility path upward from a trigger and either binding or accommodating it in one of the locations along the path, are obtained by describing all possible 'landing sites' for the presuppositional material while underspecifying where it may actually end up. Thus, as in Keller's (1997) implementation of Van der Sandt (1992) in UDRT and Asher and Lascarides's treatment in underspecified SDRT, our underspecification formalism allows us to model parts of the resolution procedure in the anaphoric account purely declaratively.

In the case of content modification, the material that is accommodated very obviously satisfies the presuppositional condition. In a way, content modification cuts short the reasoning which in our treatment justifies or explains accommodation as such. There is no explicit constraint that says that the accommodated material itself must be inferrable in the local context of the node where it is inserted. So it seems that part of the explanatory force of the presuppositional condition is lost.

This is only partially true, because for any given presupposition there is a choice to be made between satisfying it through content modification or not. Below we will propose a preference which, all other factors being equal, makes content modification more costly than global accommodation, or no accommodation at all. Each possible reading comes with a certain cost, and finding out what must be accommodated to obtain these readings, either globally or through content modification, is essential in computing the cost. Presuppositional conditions set the necessary reasoning process going. Still, as will become clear when we go through some examples, notably cases of bridging anaphora, there is some need to recompensate for lost explanatory power. This will take the form of adding constraints on background values and their local context, regulating exactly *where* presuppositional material ends up when content modification takes place.

6.4.3 Preferences

What counts as a possible landing site for presuppositional material is subject to grammatical constraints. What may subsequently be selected as preferred possible readings depends on the hearer's preferences and common sense reasoning. Various constraints and preferences have been discussed in the preceding. The anaphoric account to presupposition can be implemented to the letter by accepting local and global informativity constraints (weak or strong) and by including its preferences. The main preferences specified in Van der Sandt (1992) are, first, a preference for 'binding' over 'accommodation', and, second, a preference for

accommodation in the main DRS[45] over local or intermediate accommodation. A slightly more sophisticated version of the latter says that accommodation at the least subordinated, or 'highest' possible level allowed by the other constraints on the interpretation process is most preferred, see Geurts (2000; 2005).[46]

Given the treatment of presupposition theory in our grammar, the first preference translates as a preference for satisfying a presuppositional condition by binding the presupposition trigger to a discourse marker which was *linguistically introduced* by an explicit antecedent, over satisfying it by binding it to a discourse marker that was *accommodated* in the interpretation process. The second comes out as a preference for accommodation in B or in a background value β_k such that $\beta_k \sqsubseteq (B \oplus i_r)$ over accommodation in a background value β_k which is in the scope of one or more subordinating operators in $(B \oplus i_r)$. The preference for accommodation at the least subordinated possible level can easily be defined in terms of the subordination relation on sub-DRSs of $(B \oplus i_r)$.

In order to model the anaphoric approach of Van der Sandt (1992) as closely as possible, these preferences can be accepted as such in the framework of discourse interpretation. We may also ponder a bit about what might *explain* these preferences, and whether they could possibly be made to follow from more general principles of interpretation. In the case of the preference for binding to linguistically introduced material over binding to accommodated material, it might be argued that this is the effect of a general preference for *maximally cohesive* discourse. If binding relative to an explicit antecent results in more cohesive discourse than binding to inferred or accommodated information, the preference is explained.

But what is cohesion? Are readings resulting from accommodation, but satisfying all other constraints and preferences in the grammar, intuitively speaking less cohesive than readings resulting from plain binding? Are readings resulting from bridging anaphora perceived to be less cohesive than anaphors that resolve directly to an antecent? It seems that making this argument involves defining cohesion in terms of the preference: discourse is more cohesive if, among other factors, the anaphors in it resolve through direct binding to some explicit antecedent. This, we believe, is perfectly acceptable as a hypothesis about what establishes cohesion in discourse, but it does not add very much to what we were trying to explain.

Alternatively, the preference might be explained in terms of an *economy principle* in the sense of Kamp (2001b), on the assumption that binding an anaphor

[45]This is *global accommodation* in the sense of the anaphoric account: presuppositional material K is globally accommodated if it ends up in the main DRS of the discourse meaning, that is, if $K \sqsubseteq (B \oplus i_r)$. This can result from 'global accommodation' in our sense, but also from accommodation in any local background in i_r which is not in the scope of a negation, disjunction or implication.

[46]In fact, Geurts proposes something more general in his *Buoyency Principle*, namely that all 'backgrounded' material tends to float up to the main DRS. This applies not only to accommodated, but also to bound presuppositional material.

or a presupposition trigger to a binder that is already present in the local context due to an overt antecedent is less costly than having to infer and accommodate a binder. Kamp considers an economy principle "according to which presupposition justification must make use of as much contextual information it can get hold of, thereby minimising what remains to be supplied by accommodation." He specifies a preference which says that accommodating a DRS K is to be preferred over accommodating a DRS K' iff there is a reason for preferring K over K' but not for preferring K' over K. There is a reason for preferring K over K' if K' involves at least as many new discourse referents as K, and, given an identification of discourse markers in K' with marker in K that can be seen as playing the same argument roles, the conditions of K' entail those of K.

A requirement similar to this can be formulated in our grammar. We propose the following preference, which constrains the local background values of nodes. It may be viewed as an implementation of the economy principle.

[**Prefer smaller local backgrounds**]
For all nodes k and DRSs K, K', prefer $\beta_k = K$ over $\beta_k = K'$ if $K \sqsubseteq K'$.

The effect of this preference is that readings resulting from nontrivial content modification are predicted to be less preferred, or 'more costly', than those that do not require content modification at all. The Van der Sandtian preference for 'binding' over 'accommodation' follows from this. A further consequence is that readings in which a given anaphor is interpreted in terms of plain binding to an explicit antecedent are predicted to be preferred over readings in which the anaphor is bound via a bridging link. We could not say off-hand whether this is observationally correct, but it is open for testing.[47]

Note that the preference for smaller local backgrounds does not say anything about the costs of global accommodation or the specification of the implicit background B. As was suggested in previous sections, global accommodation comes with variable costs, depending on factors as diverse as the plausibility or implausibility of what is accommodated, the type of evidence available for what is to be accommodated, and the authority of the speaker. It may be assumed that the acceptance in B of information that counts as *common knowledge* between discussion participants according to the interpreting agent is costless. It then follows that, all other factors being equal, a reading resulting from satisfying a presupposition by binding the trigger relative to globally accommodated common knowledge is just as preferred as satisfying it through plain binding to an explicit antecedent. Correspondingly, readings that require nothing more than the acceptance of common knowledge in B are predicted to be preferred over readings that

[47]It seems to be an implicit assumption in Van Deemter and Krahmer (1998), who provide a resolution algorithm for anaphoric NP's which tries to resolve them first through plain binding; if that fails, resolution through partial matching or brdiging is attempted.

require content modification. Again this may be tested.

We haven't got an explanation yet for the other main preference in the anaphoric account, the one in favour of accommodation in the main DRS of the discourse meaning, or for accommodating 'higher up' on the accessibility path from a trigger upward. The preference for smaller background values does not make accommodation in a local background higher up on the accessibility path less costly than accommodating it lower down, in other words, the economy principle does not explain the effect. Let us consider some other factors that may possibly explain it.

It has been suggested in some places[48] that there is an overall preference for logically stronger readings over weaker ones. This would explain the preference for accommodation in the main DRS of the discourse meaning. It is not clear that the preference for logically stronger readings is observationally correct, however. As discussed in Geurts (2000), if we consider presupposition accommodation in the scope of constructions headed by a universal quantifier, the preference predicts that a presupposition triggered in the nucleus of the quantifier will preferably be accommodated in the nucleus, because the resulting reading will be stronger than the one resulting from accommodating it in the quantifier's domain. According to some people's judgements, however, the reading resulting from accommodation in the quantifier domain is preferred. Furthermore, the preference predicts that the preferred readings of quantificational sentences vary systematically with the choice of quantifier. This prediction is not born out.

Alternatively, an explanation might be given in terms of the general preference for more plausible readings. On the basis of this preference accommodation in the main DRS is to be expected whenever what is to be accommodated can be taken to be normal or uncontroversial in the actual world, or in the state of affairs the discourse is about, given the hearer's world knowledge and common sense reasoning. On the other hand, readings that describe exceptional or non-default situations given the hearer's background knowledge will be discredited. As was suggested in section 6.4.1, however, the preference for more plausible readings alone is not discriminating enough to account for all observations.

Moreover, the preference can be overruled. As is pointed out in Asher and Lascarides (1998b), there are cases in which less plausible reading are preferred over more plausible ones, while the preferred reading is one in which material is accommodated in the main DRS. Thus, a preference for plausible readings can only be a part of the explanation. The example they discuss, originating from Beaver (1997), is shown in (6.46a). The presupposition under consideration is the one triggered by 'the knowledge that David is a computer program running on a PC'. The 'intermediate reading' is paraphrased in (6.46b) the reading resulting from global accommodation in (6.46b).

[48] See Dalrymple et al. (1998), Zeevat (1999), Blutner (2000), and Geurts (2000).

(6.46) a. *If David wrote the article, then the knowledge that David is a computer program running on a PC will confound the editors.*

b. *If David wrote the article and he's a computer program running on a PC, then the knowledge that David is a computer program running on a PC will confound the editors.*

c. *David is a computer program running on a PC, and if he wrote the article, then the knowledge that David is a computer program running on a PC will confound the editors.*

The intuitively most preferred reading of (6.46a) is (6.46c). In terms of general background knowledge, however, this a less plausible reading than (6.46b). Asher and Lascarides explain this in terms of the interaction with the preference for global accommodation. The example suggests to us that it may simply be a default rule of interpretation that nominal descriptions are interpreted relative to the here and now and the actual world w_r, unless some explicit operator present in the discourse indicates otherwise, or other constraints and preferences interfere. The preference for global accommodation can be seen as an effect of this rule.

What we want to emphasize here is that preferences, by nature, can be overruled, and that in explaining the observations the grammar and the preference system must be taken as a whole. It is the interaction that explains the observations, not singular preferences. Plausibility is only one of the factors that may account for the observation that readings resulting from accommodation in the main DRS are generally preferred. We would like to mention two more factors that may contribute to an explanation.

One is the sensitivity of anaphoric or presuppositional elements to the discourse topic structure or the topicality status of their binders. In section 5.4 and section 6.2.7 it was proposed that pronouns and definite descriptions must have binders that count as topical given the topic hierarchy of the embedding discourse. When binders are accommodated in the course of the interpretation process they are subject to this requirement and must count as topical elements given the preceding discourse as well. Since the discourse topics available to a given, embedded clause may very well be established outside the scope of the operators embedding it, it is not surprising that binders, when they are accommodated, end up outside the scope of these operators as well.

The other is the potential impact of appropriateness conditions attached to specific connectives or discourse relations. For instance, it might be assumed that conditional connectives or discourse operators are most appropriately used when given the context of interpretation, world knowledge and common sense reasoning, there is a relation of logical dependence between the antecedent and the consequent, and all of the conjunctive components in the antecedent or the consequent contribute nontrivially to the dependence. Specifically it might be assumed that background material is only appropriately present in the scope of the conditional if it logically depends on the semantic value of the antecedent (and is not trivially

satisfied in the sense that the local context of the conditional already supports it). In line with our discussion in section 6.3 this could be implemented by means of soft constraints on the local contexts of the arguments of the operator. Optimally satisfying these constraints would then result in most preferred readings in which presuppositional material is only captured in the scope of a conditional if it logically depends on the antecedent, given the context of interpretation.

In sum, we have briefly considered the constraining effect of a preference for logically stronger readings, a preference for more plausible readings, requirements on the topicality status of binders, and appropriateness conditions associated with individual sentence or discourse operaters. In our view, it is not implausible that the observation that readings resulting from accommodating in the main DRS of discourse meanings are generally preferred over readings resulting from accommodating in subordinated locations can be accounted for in terms of the interaction of these constraints and preferences. This claim can only be tested in a fully developed and formally spelled out preference system, however.

Awaiting such testing, the preference for accommodating in the main DRS may be accepted as such, since it appears to be observationally correct. If the preference can be explained from principles, or independent preferences and constraints in the grammar, this only reinforces its credibility. As for the more sophisticated variant of the preference which says that accommodation at less subordation locations is preferred over accommodation at more subordinated locations, we are not certain that it is observationally correct. We feel the safest approach here is to refrain from regulating anything and trust that the subtle intuitions can eventually be explained in terms of interacting preferences and grammatical constraints.

Concluding, the result of our considerations is that we have adopted a preference for minimizing local backgrounds, which implements a form of 'economy'. The more material is accepted in local backgrounds, the more costly the corresponding reading is. We assume that in general costly readings must be avoided. The preference for 'binding' over 'accommodation' in the anaphoric approach to presupposition follows from this. The preference for 'global accommodation', that is, accommodation in the main DRS of a discourse meaning, cannot be explained in terms of economy. It was suggested, however, that this preference may well be an effect of independent preferences and constraints in framework of discourse interpretations. Awaiting further research, we adopted it explicitly in the discourse theory. Meanwhile, factors such as the preference for more plausible readings, constraints on the topical status of potential binders, consistency and informativity requirements, and appropriateness conditions associated with individual clause and discourse operators in the lexicon limit the number of possible preferred readings each in their own way.

6.5 The Extended Theory Applied

To end this chapter, let us go through some examples, and see how our 'combined approach' to presupposition theory, enriched with the possibility to modify the discourse content through the specification of local background information, fares with them. As always, the interpretation of a sentence is modeled as a reasoning process. The hearer draws conclusions about the discourse meaning on the basis of the discourse description resulting from processing the sentence in its discourse context. Employing his grammatical knowledge, world knowledge, and preferences, he determines the most preferred readings of the discourse. A substantial part of his reasoning, as we will see, concerns the acceptance of background information.

We will start with a standard case of direct binding relative to an explicit antecedent (section 6.5.1). Subsequently, we discuss a nonpresupposing conditional sentence in which informativity constraints play an important role (section 6.5.2), a presupposing conditional which motivates the introduction of 'justification conditions' on background material occurring in the scope of conditionals (section 6.5.3), a nonpresupposing conditional which illustrates the phenomenon of 'trapping' (section 6.5.4), a quantificational example which leads to some further discussion of intermediate readings and domain restriction (section 6.5.5), a case of bridging anaphora (section 6.5.6), a disjunctive nonpresupposing construction in which local consistency constraints play an essential role (section 6.5.7), and finally, a denial in dialogue, which enforces the local accommodation of a presupposition (section 6.5.8).

Since many of our examples feature possessive definite descriptions, some preliminary words on the analysis of the possessive determiner are in order. To facilitate our treatment of presupposition, we will assume that the possessive pronoun 'his' embeds an anaphoric element 'he', and adopt a syntactic analysis in which the embedded personal pronoun and the genitive morphology are separate anchors.[49] The genitive morphology introduces the main discourse marker of the DP while the pronoun generates a discourse marker of its own. The pronoun's marker fills an ownership role relative to the main marker of the DP. When the possessive is combined with a relational or functional nominal predicate, the pronoun's marker fills an internal role of the predicate. The following lexical descriptions of the possessive definite descriptions 'his car' and 'his wife' illustrate this.

[49] If this is not acceptable syntactically, alternatives can easily be provided. Suppose there is only a single terminal anchor for 'his'. It could either be assumed that a single node can come with more than one Π-value. Or it could be assumed that there is a single, maximal Π-value, and that background values are constructed from subsets of Π-values.

(6.47)

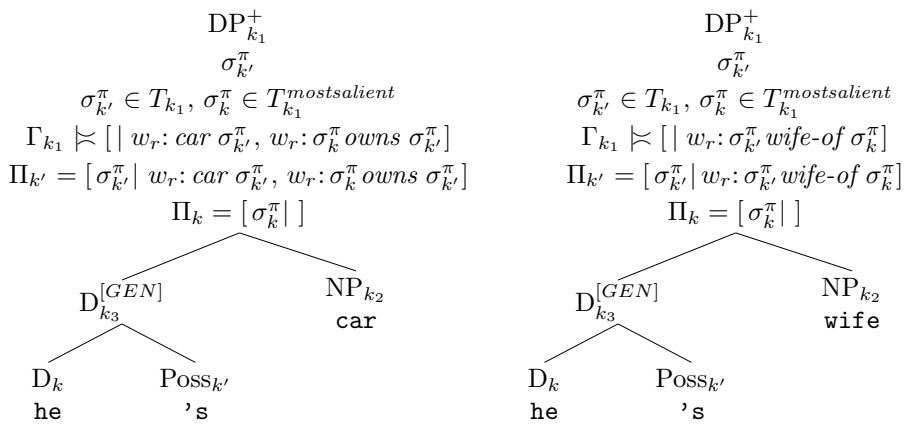

Furthermore, for convenience, we repeat the local informativity en consisteny conditions here. We will view these conditions as *preferences*, and will assume that they trigger soft constraints on the arguments of sentence and discourse operators that figure in a given discourse description. We will only represent the constraints in pictures of discourse descriptions when they are relevant to the discussion.

Local Informativity $\forall k\,[\,\ell(k)=s \rightarrow \Gamma_k \not\models \sigma_k]$

Local Consistency $\forall k\,[\,\ell(k)=s \rightarrow \Gamma_k \not\models [\,|\,\textbf{not}\ \sigma_k]\,]$

6.5.1 Binding Relative to an Overt Antecedent

Nothing is changed in our treatment of anaphora resolution and binding. As before, the Properness Requirement enforces that anaphoric elements are bound, and that presuppositional conditions are satisfied given binding of the anaphoric discourse markers in the presuppositional material. In comparison with the previous discussions, the economy principle introduced in the preceding section now leads to a more clearcut prediction that readings in which the presupposition is bound by an overt antecedent are preferred over readings that require content modification. Consider the following basic example.

(6.48) *Pedro has a mule. The animal is happy.*

After the hearer's processing of (6.48), and his reasoning about node identifications under the assumption that the discourse description is right-closed, the resulting description may be represented as follows. The DRSs at the nodes represent semantic values, as usual. We assume that discourse relation *elaboration* composes the semantic value of its topnode from the information contents of its arguments, and constructs the local contexts of its arguments just like the connective 'and', see the lexical description of 'and' in (6.43).

(6.49)

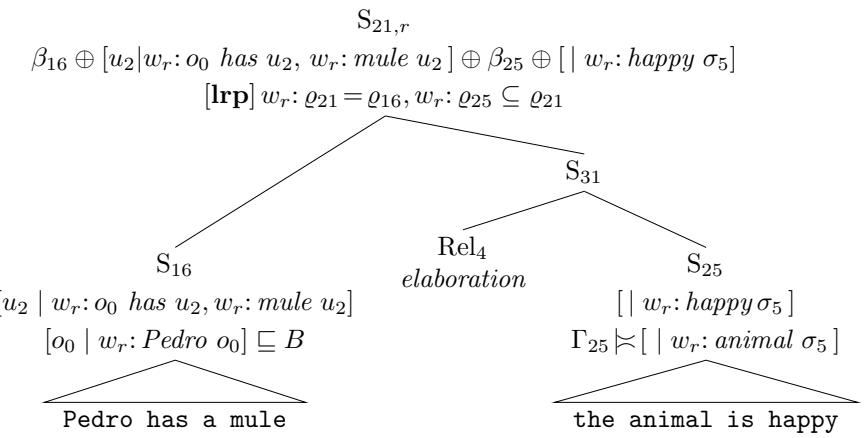

The discourse meaning $B \oplus i_r$ that can be obtained from (6.49) is given in (6.50a), where (6.50b) and (6.50c) must be satisfied. The proper name 'Pedro' introduces (6.50b), the definite description 'the animal' triggers the presuppositional condition (6.50c).

(6.50)a. $B \oplus \beta_{21} \oplus \beta_{16} \oplus [u_2 | w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2\,] \oplus \beta_{25} \oplus [\,|\, w_r\colon \mathit{happy}\ \sigma_5]$

b. $[o_0 \mid w_r\colon \mathit{Pedro}\ o_0] \sqsubseteq B$

c. $B \oplus \beta_{21} \oplus \beta_{16} \oplus [u_2 | w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2\,] \oplus \beta_{25} \mathrel{|\!\sim} [\ |\ w_r\colon \mathit{animal}\ \sigma_5\,]$

d. $[\,|\, [\,u_{22} \mid w_r\colon \mathit{mule}\ u_{22}\,] \Rightarrow [\,|\, w_r\colon \mathit{animal}\ u_{22}\,]\,]$

e. $[\,|\, w_r\colon \mathit{animal}\ o_0\,]$

There are various ways in which the presuppositional condition can be satisfied. One option is to identify $\sigma_5 = u_2$ and accept as common knowledge in B that mules are animals (6.50d). On the basis of the preference for smaller local backgrounds, all local background values can be identified with $[\,|\,]$. On the assumption that the accommodation of common knowledge is costless, this is a 'no cost' reading, and therefore a (most) preferred one.

Another option is to identify $\sigma_5 = o_0$ and accept in B that Pedro is an animal (6.50e). Again the local background values can be identified with $[\,|\,]$. As far as the preference for smaller backgrounds is concerned, this is also predicted to be a no cost reading. Intuitively, however, it is less preferred, at least when (6.48) is uttered in an out-of-the-blue context. An explanation of this may be given along the following lines. Suppose Pedro is a human being. Then it is slightly uncommon to refer to a him as an animal. Alternatively, suppose Pedro is an not a human being but an animal. Then it is somewhat uncommon to talk about an animal that owns a mule. In each case it might be argued that a somewhat exceptional or nondefault context of interpretation is presupposed. Apparently, without any preliminary priming of the context of interpretation in the utterance situation, accommodation of nondefault information is more costly than accommodation of what is perceived to be standard common knowledge. Possibly, the reason why

accommodation is more costly is that what must be accommodated has got to do with setting the parameters of the 'genre' or 'conversation type' of the discourse, rather than just accepting some world knowledge.

Finally, another set of possible solutions results from accepting presuppositional material in a local background that is part of the local context of the definite description. Given axiom $\mathcal{A}25$, there are three mutually exclusive options: $[\,\sigma_5 \mid w_r\colon \mathit{animal}\ \sigma_5\,] \sqsubseteq \beta_{21}$, or $[\,\sigma_5 \mid w_r\colon \mathit{animal}\ \sigma_5\,] \sqsubseteq \beta_{16}$ or $[\,\sigma_5 \mid w_r\colon \mathit{animal}\ \sigma_5\,] \sqsubseteq \beta_{18}$. Choosing a new discourse marker for σ_5, for instance u_5, all of these options correspond to a reading in which there is some animal, not necessarily the mule or Pedro, which is happy. On the basis of the preference for smaller local backgrounds, however, this reading is dispreferred relative to the ones resulting from binding the presupposition trigger either to 'Pedro' or to the mule. The most preferred reading of (6.48) remains the one in which 'the animal' refers to the mule Pedro owns. The corresponding discourse meaning is (6.51a) as constrained by (6.51b).

(6.51)a. $B \oplus [w_r\, o_0\, u_2 \mid w_r\colon \mathit{Pedro}\ o_0,\ w_r\colon o_0\ \mathit{has}\ u_2,\ w_r\colon \mathit{mule}\ u_2,\ w_r\colon \mathit{happy}\ u_2]$

b. $[\,|\,[\,u_{22} \mid w_r\colon \mathit{mule}\ u_{22}\,] \Rightarrow [\,|\,w_r\colon \mathit{animal}\ u_{22}\,]\,] \sqsubseteq B$

6.5.2 Nonpresupposing Conditional

Now consider an example in which the most preferred reading involves content modification. We simplify the discussion by considering only readings in which both the personal pronoun 'he' and the possessive 'his' resolve to 'John'.[50]

(6.52) *If John*[1] *is married, he*$_1$ *brings his*$_1$ *wife.*

After processing the sentence, reasoning about node identifications, and resolving both 'he' and 'his' to the most salient topical element in the context of interpretation, the hearer's discourse description may be taken to look as follows.

(6.53)

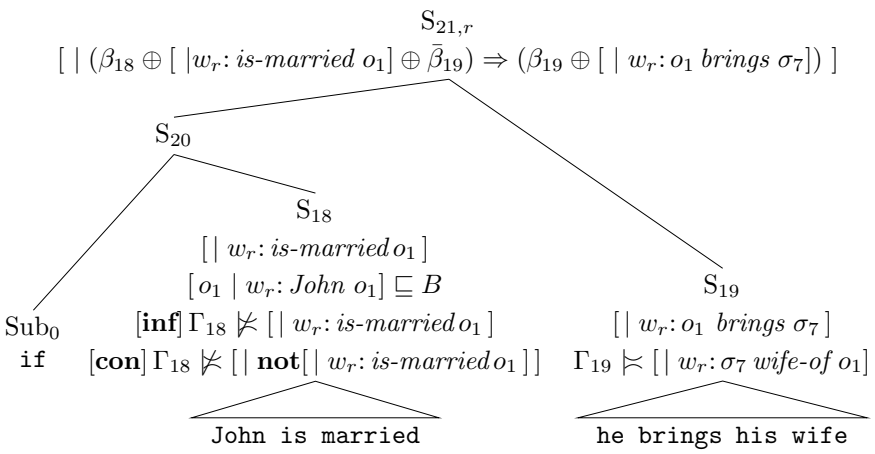

[50]The example was previously discussed in section 6.4.1, labelled as (6.35).

The discourse description defines a discourse meaning (6.54a) where the constraints in (6.54b-e) must be satisfied. The condition (6.54b) is introduced by the proper name 'John', and the definite description 'his wife' triggers the presuppositional condition (6.54c); (6.54d) and (6.54e) are the local informativity and consistency constraints associated with the antecedent of the conditional.

(6.54)a. $B \oplus \beta_{21} \oplus [|\ (\beta_{18} \oplus [\ | w_r\colon \textit{is-married}\ o_1] \oplus \bar{\beta}_{19}) \Rightarrow (\beta_{19} \oplus [\ |\ w_r\colon o_1\ \textit{brings}\ \sigma_7])]$

b. $[o_1 \mid w_r\colon \textit{John}\ o_1] \sqsubseteq B$

c. $B \oplus \beta_{21} \oplus \beta_{18} \oplus [\ |\ w_r\colon \textit{is-married}\ o_1\ ] \oplus \bar{\beta}_{19} \oplus \beta_{19} \models [\ |\ w_r\colon \sigma_7\ \textit{wife-of}\ o_1]$

d. $B \oplus \beta_{21} \oplus \beta_{18} \not\models [\ |\ w_r\colon \textit{is-married}\ o_2\ ]$

e. $B \oplus \beta_{21} \oplus \beta_{18} \not\models [\ |\ \textbf{not}[\ |\ w_r\colon \textit{is-married}\ o_1\ ]\ ]$

A simple way of satisfying the presuppositional condition would be to accept as common knowledge in B that John is married and has a wife:

$$[u_{22} \mid w_r\colon \textit{is-married}\ o_1, w_r\colon u_{22}\ \textit{wife-of}\ o_1\ ] \sqsubseteq B$$

If the hearer has no specific information about John indicating that he is not married, this is uncontroversial information and he may globally accommodate it. This would result in a cost-free reading. It can immediately be seen however, that this solution would violate the informativity and consistency conditions associated with the antecedent of the conditional, as in the anaphoric account of presupposition interpretation. Since we assume that these conditions are soft constraints, our predictions are less strong, though. The grammar predicts that the resulting reading of the discourse is dispreferred. Accommodation of the background material $[\,\sigma_7 \mid w_r\colon \sigma_7\ \textit{wife-of}\ o_1]$ in β_{21} results in a dispreferred reading for the same reason.

The other options are to accept $[\,\sigma_7 \mid w_r\colon \sigma_7\ \textit{wife-of}\ o_1]$ as background material in β_{19} or in $\bar{\beta}_{19}$; axiom $\mathcal{A}24$ excludes the inclusion of $[\,\sigma_7 \mid w_r\colon \sigma_7\ \textit{wife-of}\ o_1]$ in β_{18}. On the basis of the economy principle, the hearer will minimize local background values, and assume either $\beta_{19} = [\,\sigma_7 \mid w_r\colon \sigma_7\ \textit{wife-of}\ o_1]$ or $\bar{\beta}_{19} = [\,\sigma_7 \mid w_r\colon \sigma_7\ \textit{wife-of}\ o_1]$, while the other local background values are set to $[\,|\,]$. The first option corresponds to 'local accommodation', the second to 'intermediate accommodation'. The corresponding discourse meanings, constrained by (6.55c-d), are given in (6.55a) and (6.55b) respectively.

(6.55)a. $B \oplus [\,w_r\ o_1 \mid w_r\colon \textit{John}\ o_1$
$[\ |\ w_r\colon \textit{is-married}\ o_1] \Rightarrow [\ \sigma_7 \mid w_r\colon \sigma_7\ \textit{wife-of}\ o_1,\ w_r\colon o_1\ \textit{brings}\ \sigma_7]]$

b. $B \oplus [\,w_r\ o_1 \mid w_r\colon \textit{John}\ o_1$
$[\ \sigma_7 \mid w_r\colon \textit{is-married}\ o_1,\ w_r\colon \sigma_7\ \textit{wife-of}\ o_1\ ] \Rightarrow [\ |\ w_r\colon o_1\ \textit{brings}\ \sigma_7]]$

c. $B \not\models [\ |\ w_r\colon \textit{is-married}\ o_2\ ]$

d. $B \not\models [\ |\ \textbf{not}[\ |\ w_r\colon \textit{is-married}\ o_1\ ]\ ]$

Our treatment does not predict a difference in acceptability between these readings. The preference for accommodation in the main DRS (in B or in β_{21}) does not distinguish them. Furthermore, the topicality requirement on σ_7 is satisfied in each case by resolving the embedded anaphor to 'John', 'John' presumably being a salient element in the local topic domains of both the antecedent and the consequent of the conditional. For the example at hand it does not seem to matter, since we have no clear intuition that the one or the other reading is more preferred. We may just leave it at that and accept both readings as equally preferred.

Interestingly, the preferred readings of the sentence require content modification. In terms of the preference for selecting smaller background values, they are less than maximally preferred. As we saw, however, a violation of the preference for local informativity of the antecedent of the conditional results in a dispreferred reading. Thus, what we want our theory to say in this case is that flouting the preference for selecting smaller backgrounds is less costly than flouting the local informativity requirement, or that the second preference overrules the first. The need for further spelling out the preference system is apparent here.

6.5.3 Presupposing Conditional— Introducing Local Justification Conditions

Now compare the following example. As we saw in section 6.4.1, where (6.56) was labelled (6.36), the most preferred reading of this sentence is a presupposing one: we conclude from (6.56) that John has a wife.

(6.56) *If John^1 is coming to the conference, he$_1$ brings his$_1$ wife.*

Interpreting (6.56), the hearer will construct a discourse description analoguous to the one in (6.53). It can easily be seen that in this case, the presuppositional condition triggered by 'his wife' can be satisfied through global accommodation without violating the informativity and consistency constraints associated with the antecedent of the conditional. Accepting as common knowledge in B that John has a wife is a cost free solution and the resulting reading will be predicted to be the most preferred one. The preference for accommodating in the main DRS of a discourse meaning will strengthen this prediction.

If we want to capture the intuition that the presuppositional material is projected outside the scope of the conditional because it does not in any way depend on the antecedent of the condition, we may take up the suggestion made earlier and associate appropriateness conditions with the conditional, constraining the relation between local backgrounds and their context of interpretation. The revised lexical description of the conditional connective in (6.57) shows a way of doing this. In combination with the following general axiom, it ensures that background

material is accepted in the scope of a conditional only when it is inferrable in the local context established by the antecedent of the conditional, but not in the local context of the conditional as a whole. In other words, the background material it is *informative* outside the scope of the conditional.

$\mathcal{A}27\ \forall k[\,(k{\neq}r \wedge \beta(k){\neq}[\,|\,])\ \rightarrow\ \Gamma(top(\alpha^{-}(k)) \not\models \beta(k)\,]$

We use $top(k)$ to indicate the topnode in the elementary structure anchored by node k;[51] as before, $\alpha^{-}(k)$ is the node to which a node k is negatively anchored. For nonempty background values, axiom $\mathcal{A}27$ generates the constraints $\Gamma_{k_1} \not\models \beta_{k_2}$, $\Gamma_{k_1} \not\models \bar{\beta}_{k_3}$, and $\Gamma_{k_1} \not\models \beta_{k_3}$ in conditional structures described by (6.57).

(6.57)

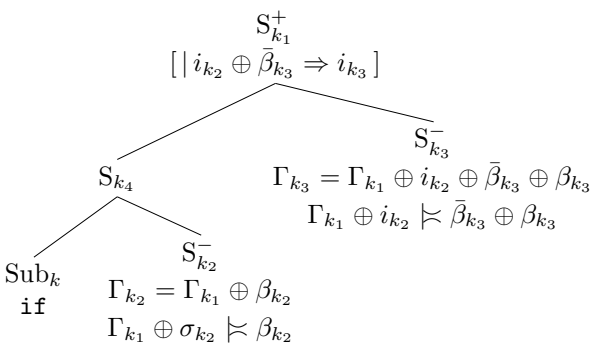

As the reader may check for himself, the combination of $\mathcal{A}27$ and (6.57) will push background material which does not depend on the antecedent of the conditional but which is satisfiable at some less embedded level in the discourse meaning outside the scope of the conditional. Note that axiom $\mathcal{A}27$ applies to all clause and discourse operators. The effect of the axiom is that we can no longer a have series of nodes upward the accessibility path from a given trigger such that accommodation at these nodes results in the same reading. Choosing a different location for accommodation will result in a different reading. If there is a series of nodes upward the accessibility path of a trigger such that in all of these positions the background material of the presupposition would be satisfiable in principle, $\mathcal{A}27$ enforces that the highest position is chosen. Thus, the axiom seems to account for the 'drift upward' of accommodated material.

Dependency constraints like the ones we associated with the conditional construction can in principle be incorporated in the lexical description of other discourse and sentence operators as well. We will generally refer to the combination of constraints triggered by $\mathcal{A}27$ plus the dependency constraints on background values associated with individual operators as *justification* conditions. They justify the insertion of background material in a certain location, when content modification takes place.

[51] $\forall k, k'[\,top(k){=}k' \leftrightarrow (k{=}\alpha^{+}(k')\ \wedge\ \neg\exists k''[k{=}\alpha^{+}(k'')\ \wedge\ k''{\lhd^{+}}k'])\,]$

6.5.4 Trapping through Direct Binding

We have seen a case in which presuppositional material was captured in the scope of a conditional because of local informativity constraints. As in the anaphoric account to presupposition, background material may also be 'trapped' in the scope of an operator because it is bound by one or more discourse markers introduced in the scope of the operator. This is what happens, for instance, in the following example.

(6.58) *If a*[1] *woman goes shopping she*$_1$ *takes her*$_1$ *cell-phone.*

We only consider readings in which the indefinite 'a woman' binds the anaphoric elements 'she' and 'her'. Suppose the hearer constructs the discourse description shown in (6.59).

(6.59)

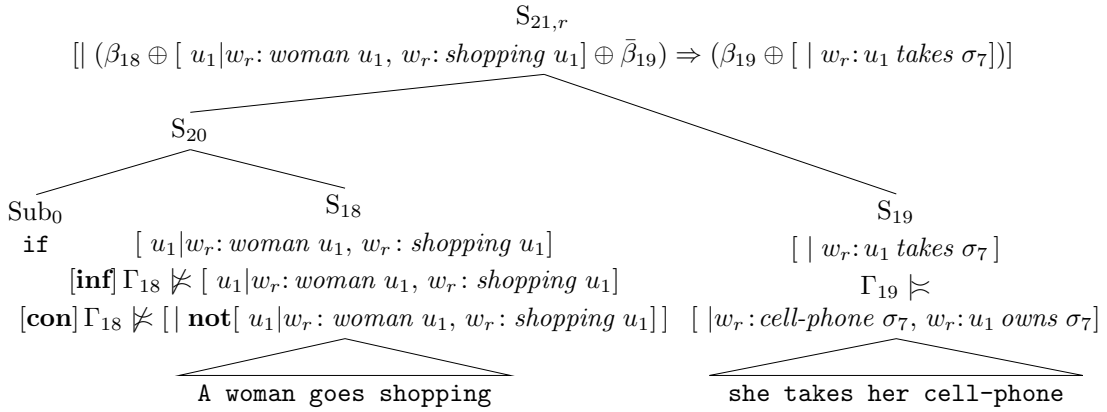

A discourse meaning of the kind (6.60a) can be computed from this, where various conditions must be satisfied. One of these is the presuppositional condition in (6.60b), which is triggered by 'her cell-phone'. Two others are the local informativity and consistency conditions associated with the antecedent of the conditional, (6.60c) and (6.60d).

(6.60)a. $B \oplus \beta_{21} \oplus [|\ (\beta_{18} \oplus [\ u_1|w_r\colon \mathit{woman}\ u_1,\ w_r\colon \mathit{shopping}\ u_1] \oplus \bar{\beta}_{19}) \Rightarrow (\beta_{19} \oplus [\ |\ w_r\colon u_1\ \mathit{takes}\ \sigma_7])]$

b. $B \oplus \beta_{21} \oplus \beta_{18} \oplus [\ u_1 \mid\ w_r\colon \mathit{woman}\ u_1,\ w_r\colon \mathit{shopping}\ u_1\ ] \oplus \bar{\beta}_{19} \oplus \beta_{19} \mathrel{|\!\sim} [\ |\ w_r\colon \mathit{cell\text{-}phone}\ \sigma_7,\ w_r\colon u_1\ \mathit{owns}\ \sigma_7]$

c. $B \oplus \beta_{21} \oplus \beta_{18} \not\models [\ u_1|w_r\colon \mathit{woman}\ u_1,\ w_r\colon \mathit{shopping}\ u_1]$

d. $B \oplus \beta_{21} \oplus \beta_{18} \not\models [\ |\ \mathbf{not}[\ u_1|w_r\colon \mathit{woman}\ u_1,\ w_r\colon \mathit{shopping}\ u_1]\ ]$

The last two conditions once more illustrate the point made in section 6.3 that the local informativity and consistency conditions, when associated as hard constraints with all discourse units, are too constraining. The conditions in (6.60c)

and (6.60d) require that the local context neither supports that there is a woman who goes shopping, nor supports that it is not the case that there is a woman who goes shopping. Given world knowledge, the second may plausibly be accepted, but the first seems unlikely. Unless we have a reason to interpret the semantic content of the antecedent of the conditional much more specifically than we do here, it is hard to satisfy the informativity condition. Intuitively speaking there is no informativity problem with (6.58), the sentence is felicitous. Whether or not there are women who go shopping in the actual world does not matter. The assumption that the local informativity conditions are soft constraints saves the account: [**inf**] can be violated without making the conditional infelicitous. What remains to be explained, however, is why violating the informativity preference is unproblematic in this case, while in, for instance, (6.52) it leads to a dispreferred reading.

Now consider the presuppositional condition in (6.60b). If we attempt to satisfy this by accommodating the information $[\sigma_7 \mid w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7]$, which might be paraphrased as "u_1 owns a cell-phone", outside the scope of the conditional, there is a problem with the Properness requirement. Discourse marker u_1 will not be bound. We assume that when a binder is accommodated a fresh discourse marker is selected. Consequently, the Properness requirement cannot be satisfied by accommodating $[\ u_1 \mid\ ]$ in addition. The only possibility is to accommodate $[\sigma_7 \mid w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7]$ in $\bar{\beta}_{19}$ or β_{19}.[52]

Suppose we accept in our discourse theory the justification conditions associated with the conditional that were introduced above. Then what are the implications for examples like (6.58)? Setting either β_{19} or $\bar{\beta}_{19}$ to $[\ |\ ]$ on the basis of the economy principle, the appropriate conditions look as follows. (6.61a) might be called the dependency constraint triggered by the conditional, (6.61b) may be referred to as the informativity requirement on local background information.

(6.61)a. $B \oplus \beta_{21} \oplus \beta_{18} \oplus [\ u_1 \mid w_r\colon \textit{woman}\ u_1,\ w_r\colon \textit{shopping}\ u_1\ ] \models [\ \sigma_7 \mid w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7]$

b. $B \oplus \beta_{21} \not\models [\ \sigma_7 \mid w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7]$

As for the informativity requirement (6.61b), note that given our definition of the entailment relation $\models$, for any pair of DRSs K, K', $K \not\models K'$ is satisfied if K' contains a free referent. Thus, (6.61b) is satisfied without having any further effect on the interpretation of the discourse. The dependency constraint (6.61a) takes on the task of the presuppositional condition at the location where content

[52]Note that, generally speaking, it is quite plausible that there is / are individuals who own cell-phones. As common knowledge, this can be accepted in B. Again, it should be avoided that in accommodating this information, the discourse marker u_1 is re-declared, resulting in requantification of the discourse marker in the output discourse meaning. In other words, we will not have $[\ u_1\, u_{22} \mid w_r : \textit{cell-phone}\ u_{22},\ w_r : u_1\ \textit{owns}\ u_{22}] \sqsubseteq B$. This can be excluded by adopting a general constraint that prohibits the occurrence of discourse markers in universes other than the one in which they are base-generated.

modification takes place. The constraint is very much like the presuppositional condition in (6.60b), the only difference is that the presuppositional material in (6.61a) is bound. (6.61a) says that the local context must entail that u_1 owns a cell-phone. Given axiom $\mathcal{A}25$ the background material $[\ \sigma_7 \mid w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7]$ can only be accommodated once.

The dependency constraint is satisfied if, for example, we are prepared to accept as common knowledge that all women own cell-phones. Alternatively, it may be assumed that the conditional is a statement about women who own cell-phones, that is, the dependency constraint induces a restriction on the elements in the topic domain. Yet another option would be to accept as common knowledge that women normally own cell-phones, and that the conditional is a statement about 'default women'. The first of the resulting readings is somewhat implausible, the other two, however, seem plausible enough. The reading in which the conditional talks about women who own cell-phones corresponds to the one resulting from intermediate accommodation (6.62a); the reading in which the conditional talks about default women corresponds to the one resulting from local accommodation (6.62b).

(6.62)a. $B \oplus$

$$[\ |\ [u_1\ \sigma_7\ |w_r\colon \textit{woman}\ u_1,\ w_r\colon \textit{shopping}\ u_1,\ w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7] \Rightarrow [\ |\ w_r\colon u_1\ \textit{takes}\ \sigma_7]\ ]$$

b. $B \oplus [\ |\ [\ u_1|w_r\colon \textit{woman}\ u_1,\ w_r\colon \textit{shopping}\ u_1] \Rightarrow [\ \sigma_7 \mid w_r\colon \textit{cell-phone}\ \sigma_7,\ w_r\colon u_1\ \textit{owns}\ \sigma_7,\ w_r\colon u_1\ \textit{takes}\ \sigma_7]\ ]$

6.5.5 Intermediate Readings

It should be clear from the above how so-called 'intermediate readings' of conditional or quantificational sentences will be treated in our grammar. Consider

(6.63) *Every*^{0} *child brings his*$_0$ *pet animal to school.*

Suppose the hearer constructs a discourse description after interpreting (6.63). He subsequently computes a discourse meaning of the kind (6.64a) from the description, where various conditions must be satisfied. One of these is the presuppositional condition shown in (6.64b), which is triggered by 'his pet animal'.

(6.64)a. $B \oplus \beta_{21} \oplus [|\ ([\ u_0|w_r\colon \textit{child}\ u_0\ ] \oplus \bar{\beta}_{19}) \Rightarrow (\beta_{19} \oplus [\ |\ w_r\colon u_0\ \textit{brings-to-school}\ \sigma_4])\ ]$

b. $B \oplus \beta_{21} \oplus [\ u_0 \mid\ w_r\colon \textit{child}\ u_0\ ] \oplus \bar{\beta}_{19} \oplus \beta_{19} \models [\ |\ w_r\colon \textit{pet-animal}\ \sigma_4,\ w_r\colon u_0\ \textit{owns}\ \sigma_4\ ]$

As with example (6.58), the presuppositional condition cannot be satisfied by accepting as common knowledge that all children own pet-animals, and accepting the information $[\ \sigma_4 \mid w_r\colon \textit{pet-animal}\ \sigma_4,\ w_r\colon u_0\ \textit{owns}\ \sigma_4\ ]$ either in B or in

β_{21}. Discourse marker u_0 would not be bound and the Properness Condition on (6.64a) prohibits this. The background material can only be accepted in $\bar{\beta}_{19}$ or β_{19}, resulting in the 'intermediate' and the 'local' reading, respectively.

Just like the conditional connective, the quantifier 'every' can be decorated with conditions on background values in its scope, which ensure that the insertion of background material in the restrictor or the nucleus of the quantificational structure is justified given the local context. We may adopt the following lexical description of 'every'. Axiom $\mathcal{A}27$ will trigger $\Gamma_{k_1} \not\models \bar{\beta}_{k_2}$ and $\Gamma_{k_1} \not\models \beta_{k_2}$ for nonempty $\bar{\beta}_{k_2}$, β_{k_2}.

(6.65)

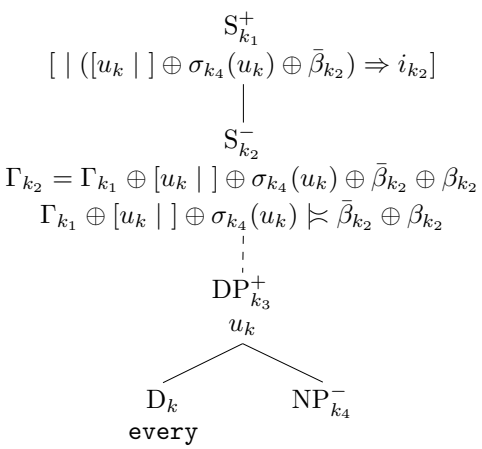

This results in predictions analoguous to the ones we had in the case of (6.58). The local dependency constraint $\Gamma_{k_1} \oplus [u_k\ |\ ] \oplus \sigma_{k_4}(u_k) \models \bar{\beta}_{k_2} \oplus \beta_{k_2}$ will be satisfied if we are prepared to accept as common knowledge that every child has a pet animal. This is implausible, however, because it concerns all children in the domain of discourse. Alternatively, it may be assumed that (6.63) is a statement about children that own pet-animals, that is, the dependency constraint induces a restriction on the elements in the topic domain. This reading may be taken to correspond to the one resulting from intermediate accommodation, (6.66a). Yet another option would be to assume that children normally have pet-animals, and that the conditional is a statement about 'default children'. This can be taken to correspond to the reading resulting from local accommodation, (6.66b).

(6.66)a. $B \oplus [\ |\ [\ u_0\, \sigma_4 \,|\, w_r\text{: } \textit{child}\ u_0,\ w_r\text{: } \textit{pet-animal}\ \sigma_4,\ w_r\text{: } u_0\ \textit{owns}\ \sigma_4] \Rightarrow$
$[\ |\ w_r\text{: } u_0\ \textit{brings-to-school}\ \sigma_4]\ ]$

b. $B \oplus [\ |\ [\ u_0 \,|\, w_r\text{: } \textit{child}\ u_0\ ] \Rightarrow$
$[\ \sigma_4\ |\ w_r\text{: } \textit{pet-animal}\ \sigma_4,\ w_r\text{: } u_0\ \textit{owns}\ \sigma_4,\ w_r\text{: } u_0\ \textit{brings-to-school}\ \sigma_4]\ ]$

We leave open how domain restriction should be implemented, and exactly how it relates to the possibility of accommodating in subsidiary background values in our grammar. Technically speaking, we can envisage alternative ways in which 'intermediate readings' can be accounted for. For instance, subsidiary background

values might be eliminated so that only standard background values would be available for the insertion of background material. In addition, it might be assumed that quantifiers link up anaphorically to a domain of elements established *outside* the scope of the quantificational structure. That domain may be partially underspecified. Subject to various other constraints, it may be strengthened as a result of the interpretation process as a whole. Intermediate readings could then be accounted for in terms of a combination of local accommodation and domain restriction of the quantifier domain.

A full-fledged account of domain restriction with quantifiers must explain under what conditions domain restriction, as induced by presupposition, is possible, and when it is not. For instance, it should somehow account for the contrast between the following examples.

(6.67)a.# *Few of the team members can drive, but every team member will come to the match in her car.*

b. *Few of the team members can drive, but every team member who owns a car will come to the match in her car.*

Geurts and van der Sandt (1998), ex. 34, based on Beaver (1994).

Intuitively, the statement in (6.67a) is infelicitous, because it suggests that every team member owns a car. Given world knowledge and common sense reasoning, this is incompatible with the information that few team members own a car. In (6.67b), there is no problem, because the second conjunct explicitly talks only about team members who own a car. As it stands, (6.67a) presents a problem for our account, because nothing prevents the intermediate accommodation of the presupposition "x owns a car" resulting in domain restriction of the quantifier in the second conjunct. The theory predicts that (6.67b) is a possible reading of (6.67a). Following Geurts and van der Sandt (1998) it may be assumed that intermediate accommodation, and the restriction of the quantifier domain, is excluded because the domain is already bound to the total set of team members. It must be assumed that this binding of the quantifier domain takes priority over the restriction of that domain that is potentially induced by presupposition interpretation. Note that the binding of 'every team member' in (6.67a) to the superset of 'few team members' is reinforced by the relation of contrastive parallelism that links the arguments of 'but'.

6.5.6 Bridging

Bridging anaphors relate to an antecedent in the preceding discourse through an implicit bridging link. Although the link is implicit one would expect bridging anaphors to behave exactly like elements that contain an overt embedded anaphor, such as, for instance, possessive definite descriptions. Furthermore, the

dependency and informativity constraints associated with conditionals can be expected to come in useful, in that they will help to compute or abduce a locally relevant bridging link, when content modification occurs. As will be shown, this is correct, but only if we are willing to adopt an analysis of definite descriptions in which the underspecified bridging link is part of the presuppositional material.

Consider (6.68), which is a variant of example (6.19), discussed in section 6.2.7. We concentrate on readings where 'he' resolves to 'John'.

(6.68) *If John[1] is married, he$_1$ blackmailed the priest.*

Intuitively, in interpreting (6.68) we assume that there was a priest involved in John's marriage in some way and we resolve the priest whom John blackmailed to that individual. This still allows for different readings. The most preferred reading, in our view, is one resulting from local accommodation, which can be paraphrased as "If John is married then there is an arbitrary priest involved in his marriage and John blackmailed him." There is also a possible reading resulting from global accommodation, which may be paraphrased as "There is a certain priest, and if John is married, he blackmailed him." With additional strengthening of the implicit background B with information about this priest and John's relation to him[53] this is an acceptable reading as well, but we will focus on the reading resulting from local accommodation here.

Interpreting (6.68), the hearer may construct the following discourse description.

(6.69)

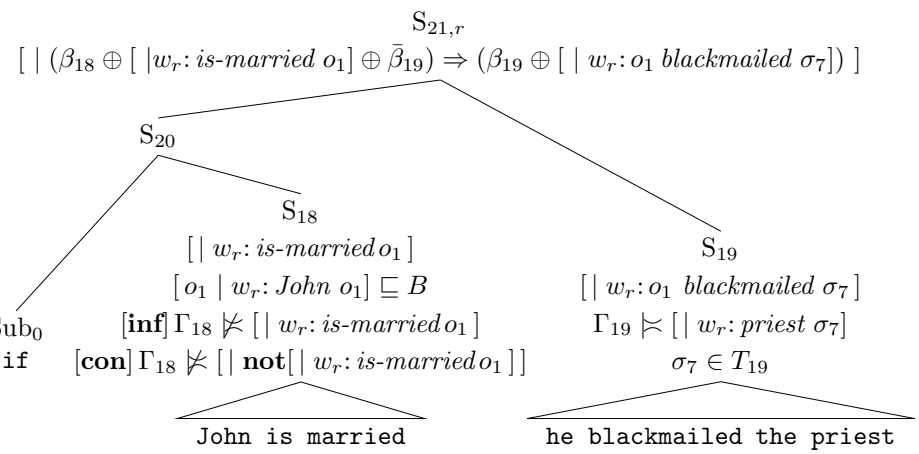

He may compute discourse meaning (6.70a) from this, where various conditions must be satisfied. One of these is the presuppositional condition shown in (6.70b), triggered by 'the priest'. The other (6.70c) is the local informativity constraint associated with the antecedent of the conditional.

[53]For instance, he is the priest performing all marriages in John's neighbourhood, and known to be very strict on certain issues, and that's why John needed to blackmail him.

(6.70)a. $B \oplus \beta_{21} \oplus$
$[\ |\ (\beta_{18} \oplus [\ |w_r\text{: } \textit{is-married}\ o_1\ ] \oplus \bar{\beta}_{19}) \Rightarrow (\beta_{19} \oplus [\ |\ w_r\text{: } o_1\ \textit{blackmailed}\ \sigma_7])\ ]$
b. $B \oplus \beta_{21} \oplus \beta_{18} \oplus [\ |\ w_r\text{: } \textit{is-married}\ o_1,\ ] \oplus \bar{\beta}_{19} \oplus \beta_{19} \models [\ |\ w_r\text{: } \textit{priest}\ \sigma_7\ ]$
c. $[\ u_{22}\ |\ w_r\text{: } \textit{priest}\ u_{22},\ w_r\text{: } u_{22}\ \textit{married}\ o_1\ ]$
d. $B \oplus \beta_{21} \not\models [\ |w_r\text{: } \textit{is-married}\ o_1\ ]$

To satisfy the presuppositional condition, B may be strengthened with suitable background information—as in our discussion of (6.19)—on the basis of which it may be inferred that there is a priest, namely the priest who performed John's marriage. The inferred information may look like (6.70c). The informativity requirement (6.70d) prevents that this is accepted in B or in β_{21}. It follows that the presuppositional condition is satisfied through content modification. If we work this out, however, there is a surprise: in satisfying the presuppositional condition through content modification, the bridging link is lost. The background material introduced by the definite description 'the priest' is $[\ \sigma_7\ |\ w_r\text{: } \textit{priest}\ \sigma_7\ ]$. Accepting this material in either β_{19} or $\bar{\beta}_{19}$ and minimizing local background values results in discourse meanings (6.71a) and (6.71b), constrained by various other conditions.

(6.71)a. $B \oplus [\ |\ \ [\ |\ w_r\text{: } \textit{is-married}\ o_1\ ] \Rightarrow [\ \sigma_7\ |\ w_r\text{: } \textit{priest}\ \sigma_7,\ w_r\text{: } o_1\ \textit{blackmailed}\ \sigma_7]\ ]$
b. $B \oplus [\ |\ [\ \sigma_7\ |\ w_r\text{: } \textit{priest}\ \sigma_7,\ w_r\text{: } \textit{is-married}\ o_1\ ] \Rightarrow [\ |\ w_r\text{: } o_1\ \textit{blackmailed}\ \sigma_7]\ ]$

As in the anaphoric approach to presupposition projection, the descriptive material plus a binder is inserted in an accessible location in the discourse meaning. We could leave it at this, and assume that on the basis of the preference for plausible readings, the hearer selects an implicit background B which explains why there is a priest if John is married. But in fact, this is exactly what the justification conditions associated with the conditional were supposed to regulate. Do they not enforce the compution of the bridging link? In the example under consideration, the dependency constraint and the informativity constraint triggered by the conditional are instantiated as follows.

(6.72)a. $B \oplus \beta_{21} \oplus [\ |w_r\text{: } \textit{is-married}\ o_1\ ] \models [\ \sigma_7\ |\ w_r\text{: } \textit{priest}\ \sigma_7\ ]$
b. $B \oplus \beta_{21} \not\models [\ \sigma_7\ |\ w_r\text{: } \textit{priest}\ \sigma_7\ ]$

Since common knowledge tells us that there are priests, (6.72a) is easily satisfied. But then (6.72b) cannot be satisfied. Even when we take into account that σ_7 must be in the local topic domain, and allow inferrable entities such as 'the priest who married John' in the topic domain, it is not clear that this would result in restricting the domain of individuals relative to which $[\ \sigma_7\ |\ w_r\text{: } \textit{priest}\ \sigma_7\ ]$ is evaluated. The safest approach to the problem seems to be to represent the potential underspecified bridging link in the descriptive material of the definite description. In section 6.2.7 a proposal to this effect was made. The extended analysis of the definite description is shown in (6.73).

(6.73) *definite det:*

$$\begin{array}{c} \mathrm{DP}^{+}_{k_1} \\ \sigma^{\pi}_{k} \\ \sigma^{\pi}_{k_2} \in T_{k_1} \\ \Gamma_{k_1} \models \sigma^{(\pi\tau)}_{k_2}(\sigma^{\pi}_{k}) \oplus R^{\pi(\pi\tau)}(\sigma^{\pi}_{k_2})(\sigma^{\pi}_{k}) \\ \Pi_k = [\sigma^{\pi}_{k} \,|\,] \oplus \sigma^{(\pi\tau)}_{k_2}(\sigma^{\pi}_{k}) \oplus R^{\pi(\pi\tau)}(\sigma^{\pi}_{k_2})(\sigma^{\pi}_{k}) \\ \diagup \quad \diagdown \\ \mathrm{D}_k \qquad \mathrm{NP}^{-}_{k_2} \\ \texttt{the} \end{array}$$

Suppose we employ the lexical description in (6.73). Now, if we locally accommodate the background material carried by 'the priest' in the example under consideration, the resulting discourse meaning will feature the underspecified bridging relation R. Choosing 'John' as a potential antecedent, the justification constraints on background values associated with the conditional look as follows.

(6.74)a. $B \oplus \beta_{21} \oplus [\ |w_r\colon \textit{is-married}\ o_1\ ] \models [\ \sigma_7 \mid w_r\colon \textit{priest}\ \sigma_7\ ] \oplus R^{\pi(\pi\tau)}(o_1)(\sigma_7)$

b. $B \oplus \beta_{21} \not\models [\ \sigma_7 \mid w_r\colon \textit{priest}\ \sigma_7\ ] \oplus R^{\pi(\pi\tau)}(o_1)(\sigma_7)$

The constraints in (6.74) are satisfied if, for instance, $R = \lambda v.w[\ |\ w_r : v\ \textit{married}\ w\ ]$, and common knowledge that supports this inference, such as discussed in section 6.2.7), is accepted in B. Thus, in this set-up, the dependency and the informativity condition associated with the conditional connective induce the abduction of the bridging link. At the same time they justify the accommodation of the background material in the scope of the conditional. The discourse meanings representing the two resulting, nonpresupposing readings are shown in (6.75a) and (6.75b), as constrained by further conditions.

(6.75)a. $B \oplus [\ |\ \ [\ |\, w_r\colon \textit{is-married}\ o_1\ ] \Rightarrow$
$[\ \sigma_7 \,|\, w_r\colon \textit{priest}\ \sigma_7,\ w_r : \sigma_7\ \textit{married}\ o_1,\ w_r\colon o_1\ \textit{blackmailed}\ \sigma_7]\ ]$

b. $B \oplus [\ |\ [\ \sigma_7 \,|\, w_r\colon \textit{priest}\ \sigma_7,\ w_r : \sigma_7\ \textit{married}\ o_1,\ w_r\colon \textit{is-married}\ o_1\ ] \Rightarrow$
$[\ |\ w_r\colon o_1\ \textit{blackmailed}\ \sigma_7]\ ]$

We leave open what route to follow from here. On the one hand, the anaphoric approach to presupposition interpretation of Van der Sandt (1992) can be modeled in our grammar without adopting the additional appropriateness constraints associated with conditionals here, and it simply does not make any claims about bridging anaphora. On the other hand, it is obviously desirable to extend the theory to cover bridging anaphors. They are quite common, and certainly a theory of presupposition that combines anaphoricity with inference must have something to say about them. Thus, it will be worthwile to work out the analysis proposed here.

A central question to be answered shall be whether the format presented above can be employed for *all* definite descriptions. How does the analysis work out for

definite descriptions which resolve to an explicit antecedent directly? Are there cases in which there really is *no* anaphoric link of any kind, in which the definite description is still felicitously used?[54] And what about possesive descriptions? Is the built-in ownership relation the bridging link by definition, or is there an independent underspecified bridging link in addition? Furthermore, given the important role of the justification conditions on local backbground values, it will be necessary to specify them for all clause and discourse operators in the grammar.

6.5.7 Nonpresupposing Disjunction

In the preceding examples, the preference for nonpresupposing readings was explained in terms of the interaction with local informativity constraints and properness requirements. Local *consistency* constraints play a role in accounting for nonpresupposing readings of disjunctive constructions. Consider the following sentence, which is a variant of (6.38a) discussed in section 6.4.1. We concentrate on readings in which the possessive, and the pronoun 'her' resolve to 'Wilma'.

(6.76) *Either Wilma*[1] *is not married or her*$_1$ *husband left her*$_1$ *long ago.*

After processing (6.76), the hearer may compose a discourse description of the following kind.

(6.77)

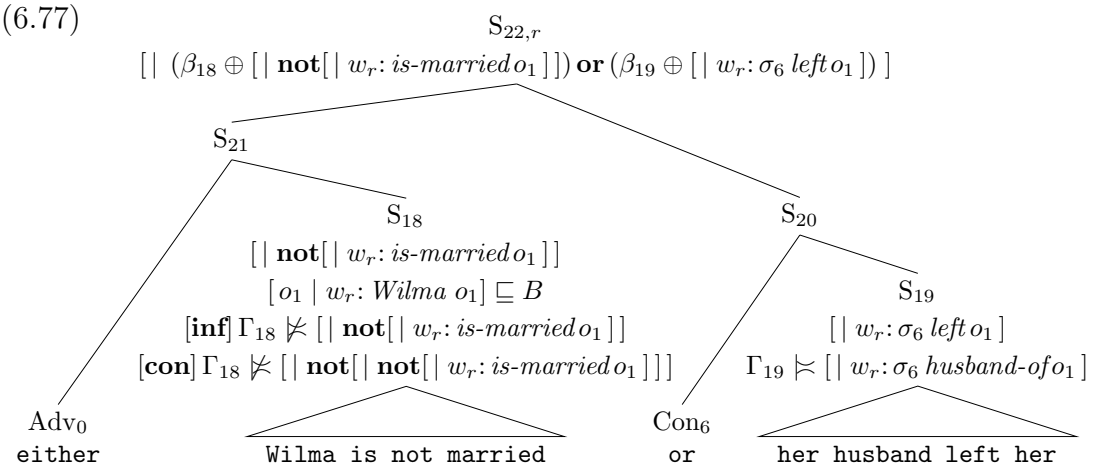

The hearer may compute discourse meaning (6.78a) from this, as constrained by various conditions. Among these are the presuppositional condition triggered by 'her husband' (6.78b), the informativity constraint associated with the first disjunct (6.78c), and the consistency requirement associated with the the first disjunct (6.78d).

[54] Heim's (1982) analysis of definites in terms of familiarity implies that they are always anaphoric in some way (see in particular the discussion on pp. 364–384). Note that accommodation is a way of satisfying the familiarity requirement.

(6.78)a. $B \oplus \beta_{22} \oplus [\,|\ (\beta_{18} \oplus [\,|\ \mathbf{not}[\,|\ w_r\colon \textit{is-married}\, o_1\,]\,])\ \mathbf{or}\ (\beta_{19} \oplus [\,|\ w_r\colon \sigma_6\ \textit{left}\, o_1\,])\,]$

b. $B \oplus \beta_{22} \oplus [\,|\ \mathbf{not}(\beta_{18} \oplus [\,|\ \mathbf{not}[\,|\ w_r\colon \textit{is-married}\, o_1\,]\,])\,] \oplus \beta_{19} \mid\!\sim [\,|\ w_r\colon \sigma_6\ \textit{husband-of}\, o_1\,]$

c. $B \oplus \beta_{22} \oplus \beta_{18} \not\mid\!\sim [\,|\ \mathbf{not}[\,|\ w_r\colon \textit{is-married}\, o_1\,]\,]$

d. $B \oplus \beta_{22} \oplus \beta_{18} \not\mid\!\sim [\,|\ \mathbf{not}[\,|\ \mathbf{not}[\,|\ w_r\colon \textit{is-married}\, o_1\,]\,]\,]$

Consider (6.78b). The presuppositional condition can in principle be satisfied by accepting as common knowledge in B that Wilma has a husband. Alternatively, this information can be accepted in β_{22} or in β_{19}. Note that acceptance in β_{18} is prevented by axiom $\mathcal{A}24$, moreover, it would not help to satisfy (6.78b). Of the three options that are still open, the interaction with the consistency requirement (6.78d) excludes accepting $[\,\sigma_6 \mid w_r\colon \sigma_6\ \textit{husband-of}\, o_1\,] \sqsubseteq B$ or $[\,\sigma_6 \mid w_r\colon \sigma_6\ \textit{husband-of}\, o_1\,] \sqsubseteq \beta_{22}$. Note that it follows from (6.78d) that $B \oplus \beta_{22} \oplus \beta_{18} \not\mid\!\sim [\,|\ w_r\colon \textit{is-married}\, o_1\,]$.[55] Since common knowledge tels us that when a woman has a husband, she is married, it is clear that this condition cannot be reconciled with accommodating $[\,\sigma_6 \mid w_r\colon \sigma_6\ \textit{husband-of}\, o_1\,]$ in B or in β_{22}. The only option is to accept $[\,\sigma_6 \mid w_r\colon \sigma_6\ \textit{husband-of}\, o_1\,] \sqsubseteq \beta_{19}$. Given minimization of local background values, this results in the following discourse meaning, as constrained by various other conditions:

(6.79) $B \oplus [\,w_r\ o_1 \mid w_r\colon \textit{Wilma}\ o_1,$
$[\,|\ \mathbf{not}[\,|\ w_r\colon \textit{is-married}\, o_1\,]\,]\ \mathbf{or}\ [\,\sigma_6 \mid w_r\colon \sigma_6\ \textit{husband-of}\, o_1,\ w_r\colon \sigma_6\ \textit{left}\, o_1\,]\,]$

Note that Global Consistency, which requires discourse meanings to be consistent, does not enforce local accommodation. If we would accommodate the background material outside the scope of the disjunction, this would not result in an inconsistent discourse meaning. In the corresponding reading, the first disjunct would be untrue in the context of interpretation and hence irrelevant, however. The local consistency requirements associated with the disjuncts prevent this.

Although there is a lot of room for variation in the analysis of disjunctive constructions, as far as we can see the justification conditions that may be associated with the disjunctive construction will be satisfied in the nonpresupposing reading. Axiom $\mathcal{A}27$, which requires the informativity of local background values outside the scope of the operator that anchors them, triggers the justification constraint $\Gamma_{19} \not\mid\!\sim \beta_{19}$ for the second conjunct. This would be instantiated as $B \oplus \beta_{22} \not\mid\!\sim [\,\sigma_6 \mid w_r\colon \sigma_6\ \textit{husband-of}\, o_1\,]$. If B contains the common knowledge that only married woman have husbands, then if the informativity constraint on the first disjunct (6.78c) is satisfied, this requirement must be satisfied as well. As for dependency conditions on the background values of the arguments of the disjunc-

[55] In general we have, for a pair of DRSs Γ, K such that $\Gamma \oplus K$ is proper, if $\Gamma \not\mid\!\sim [\,|\ \mathbf{not}[\,|\ \mathbf{not} K\,]\,]$ then $\Gamma \not\mid\!\sim K$.

tion, given our asymmetric analysis of disjunction, we might consider assigning one to the second disjunct. We might have, for example

(6.80)

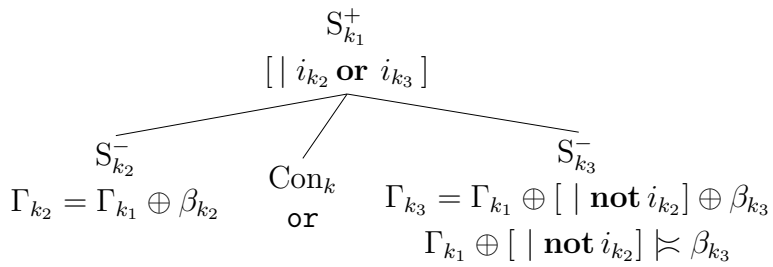

Instantiating the dependency condition in the example under discussion will produce the requirement

(6.81) $B \oplus \beta_{22} \oplus [\,|\,$ **not**$[\,|\, \beta_{18} \oplus [\,|\,$ **not**$[\,|\, w_r\colon \textit{is-married}\, o_1\,]\,]\,] \mathrel{\models\!\!\sim} \beta_{19}$

The hearer may now reason as follows. Suppose $[\,\sigma_6 \mid w_r\colon \sigma_6\, \textit{husband-of}\, o_1\,] \sqsubseteq \beta_{19}$. It then follows from (6.81) that the following condition must be satisfied

(6.82) $B \oplus \beta_{22} \oplus [\,|\,$ **not**$[\,|\, \beta_{18} \oplus [\,|\,$ **not**$[\,|\, w_r\colon \textit{is-married}\, o_1\,]\,]\,] \mathrel{\models\!\!\sim} [\,\sigma_6 \mid w_r\colon \sigma_6\, \textit{husband-of}\, o_1\,]$

On the basis of economy, set $\beta_{22} = [\,|\,]$ and $\beta_{18} = [\,|\,]$. It then follows the condition reduces to

(6.83) $B \oplus [\,|\,$ **not**$[\,|\,$ **not**$[\,|\, w_r\colon \textit{is-married}\, o_1\,]\,]\,] \mathrel{\models\!\!\sim} [\,\sigma_6 \mid w_r\colon \sigma_6\, \textit{husband-of}\, o_1\,]$

On the assumption that B contains the common knowledge that when a woman is married she has a husband, and that Wilma is a woman, it holds that

(6.84) $B \oplus [\,|\,$ **not**$[\,|\,$ **not**$[\,|\, w_r\colon \textit{is-married}\, o_1\,]\,]\,] \mathrel{\models\!\!\sim} [\,|\, [\,|\, w_r\colon \textit{is-married}\, o_1\,] \Rightarrow [\,\sigma_6 \mid w_r\colon \sigma_6\, \textit{husband-of}\, o_1\,]\,]$

It also holds that

(6.85) $B \oplus [\,|\,$ **not**$[\,|\,$ **not**$[\,|\, w_r\colon \textit{is-married}\, o_1\,]\,]\,] \mathrel{\models\!\!\sim} [\,|\, w_r\colon \textit{is-married}\, o_1\,]$

The hearer may conclude from (6.84) and (6.85) that

(6.86) $B \oplus [\,|\,$ **not**$[\,|\,$ **not**$[\,|\, w_r\colon \textit{is-married}\, o_1\,]\,]\,] \mathrel{\models\!\!\sim} [\,\sigma_6 \mid w_r\colon \sigma_6\, \textit{husband-of}\, o_1\,]$

which shows that, given reasonable assumptions about common knowledge in the implicit background B, the requirement (6.83) holds. In other words, the background value in the second disjunct is justified in the local context established by the first conjunct.

6.5.8 Denial

We like to end this section with a rather different kind of example than is usually discussed when local accommodation is concerned, namely cases of presupposition denial or rejection at discourse level. The following exchange, which is a simplified version of (6.22) discussed in section 6.3, illustrates the phenomenon.

(6.87) A: *John made an appointment with Mark's secretary.*
B: *Mark doesn't have a secretary.*

Participant A's utterance triggers the presupposition that Mark has a secretary. This is subsequently contradicted by participant B. We shall argue that in the course of the discourse, the presupposition is first accommodated in the main DRS, but next, as a result of participant B's reply, accommodated locally. The presuppositional material is then captured in the scope of a negation introduced by discourse relation *denial*.

Let us place ourselves in the position of participant B. After processing A's utterance, he may construct a discourse description of the following kind.

(6.88) Δ_1^{rcl}

$$S_{16,r}$$
$$[\,|\, w_r{:}\, o_0 \;\textit{made-app-with}\; \sigma_6\,]$$
$$[o_0 \mid w_r{:}\, \textit{John}\; o_0] \sqsubseteq B$$
$$[o_5 \mid w_r{:}\, \textit{Mark}\; o_5] \sqsubseteq B$$
$$\Gamma_{16} \mathrel{|\!\sim} [\,|\, w_r{:}\, \textit{secretary}\; \sigma_6,\; w_r{:}\, o_5\;\textit{owns}\;\sigma_6\,]$$
$$[\mathbf{inf}]\,\Gamma_{16} \not\mathrel{|\!\sim} [\,|\, w_r{:}\, o_0\;\textit{made-app-with}\;\sigma_6\,]$$
$$[\mathbf{con}]\,\Gamma_{16} \not\mathrel{|\!\sim} [\,|\, \mathbf{not}[\,|\, w_r{:}\, o_0\;\textit{made-app-with}\;\sigma_6\,]\,]$$

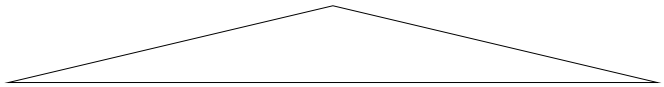

`John made an appointment with Mark's secretary`

A discourse meaning can be inferred from this in the usual way. With some reasoning, participant B may deduce discourse meaning (6.89a), where the constraints in (6.89b-d) must be satisfied. (6.89b) is the presuppositional condition triggered by 'Mark's secretary', (6.89c) and (6.89d) are the constraints triggered by the local informativity and the local consistency conditions, respectively.

(6.89)a. $B \oplus \beta_{16} \oplus [\, w_r\, o_0\, o_5 \mid w_r{:}\, \textit{John}\; o_0,\; w_r{:}\, \textit{Mark}\; o_5,\; w_r{:}\, o_0\;\textit{made-app-with}\;\sigma_6\,]$

b. $B \oplus \beta_{16} \mathrel{|\!\sim} [\,|\, w_r{:}\, \textit{secretary}\;\sigma_6,\; w_r{:}\, o_5\;\textit{owns}\;\sigma_6\,]$

c. $B \oplus \beta_{16} \not\mathrel{|\!\sim} [\,|\, w_r{:}\, o_0\;\textit{made-app-with}\;\sigma_6\,]$

d. $B \oplus \beta_{16} \not\mathrel{|\!\sim} [\,|\, \mathbf{not}[\,|\, w_r{:}\, o_0\;\textit{made-app-with}\;\sigma_6\,]\,]$

Consider the presuppositional condition (6.89b). This can only be satisfied by accepting as background information that there is an individual who is Mark's secretary. As participant B's reply shows, this is in fact incompatible with his

private knowledge or beliefs. Accommodating the presupposition will violate B's preference for plausible readings. Nonetheless we assume that, for the sake of getting A's meaning, participant B accommodates that there is an individual who is Mark's secretary. In principle there are two options. He may accept either $[\,\sigma_6 \mid w_r\colon \mathit{secretary}\,\sigma_6,\ w_r\colon o_5\ \mathit{owns}\,\sigma_6\,] \sqsubseteq B$ or $[\,\sigma_6 \mid w_r\colon \mathit{secretary}\,\sigma_6,\ w_r\colon o_5\ \mathit{owns}\,\sigma_6\,] \sqsubseteq \beta_{16}$.

Each of these results in a discourse meaning which entails that there is an individual who is Mark's secretary. Since we have been assuming all along that what is accepted in the general background B counts as common knowledge, which is *shared* between the participants in the conversation, the first option implies assigning the status of common knowledge to the presupposition. Moreover, if we maintain the semantics specified in section 4.4.7 (but compare the alternative in section 6.2.4) discourse markers declared in the universe of B must be o-markers, and will get a referential interpretation. This means that 'Mark's secretary' would get a referential interpretation. If the background material is accepted in β_{16} it does not receive the status of common knowledge and the definite description can get an existential interpretation. Depending on whether participant B wants to conform to what he thinks A intends to convey, or to his private knowledge, he may prefer either the first or the second option. We assume that both are possible readings.

The informativity and consistency conditions in (6.89c) and (6.89d) prevent that any information is accepted in B from which it would follow that John does (or does not) have an appointment with Mark's secretary.

Now participant B utters his reply. We suppose that in the preceding stage of the discourse, participant A, or an arbitrary interpreting agent, built up a discourse description analoguous to the one represented in (6.88). Processing B's contribution, he may construct an updated description of the following kind.

(6.90) Δ_2^{rcl}

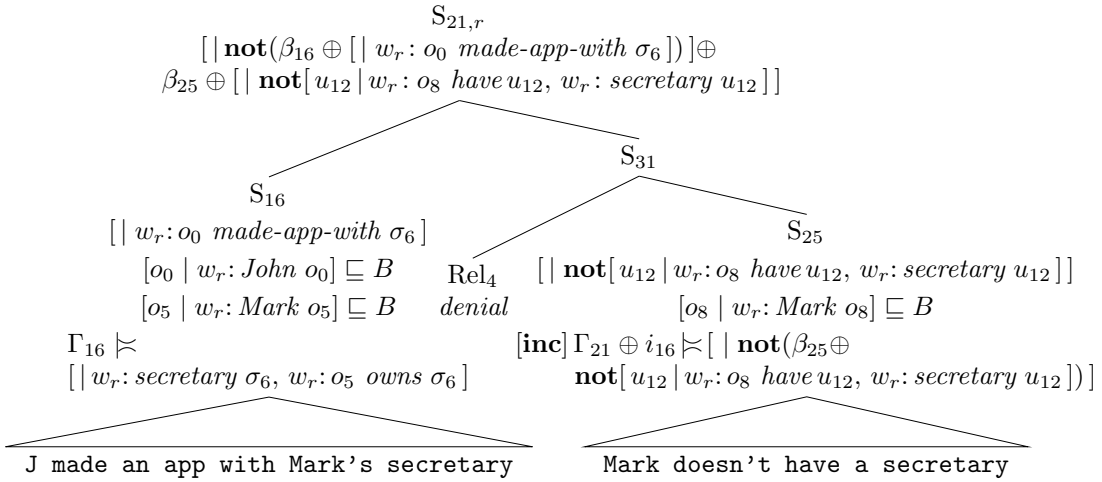

A new discourse meaning (6.91a) can be computed from this, where several conditions must be satisfied. Some are triggered by the names occurring in the discourse (6.91b-d), others are the presuppositional condition that we had before (6.91e) and the *incompatibility requirement* [**inc**] introduced by the *denial* relation (6.91f). Incompatibility is a characterising condition of denials which enforces that the semantic content of the denial is incompatible with the information content of the statement it denies.[56] Thus, there is a reason for rejecting or negating the information that is denied. The incompatibility requirement overrules or makes superfluous the local informativity and consistency constraints that would apply to node 25. The informativity and consistency constraints on node 16 (6.91g) and (6.91h) still apply, however.

(6.91)a. $B \oplus \beta_{21} \oplus [\,|\, \mathbf{not}(\beta_{16} \oplus [\,|\, w_r\colon o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,])\,] \oplus$
$\beta_{25} \oplus [\,|\, \mathbf{not}[\,u_{12} \,|\, w_r\colon o_8\ \mathit{have}\, u_{12},\ w_r\colon \mathit{secretary}\ u_{12}\,]\,]$

b. $[o_0 \mid w_r\colon \mathit{John}\ o_0] \sqsubseteq B$

c. $[o_5 \mid w_r\colon \mathit{Mark}\ o_5] \sqsubseteq B$

d. $[o_8 \mid w_r\colon \mathit{Mark}\ o_8] \sqsubseteq B$

e. $B \oplus \beta_{21} \oplus \beta_{16} \mathrel{\models\!\!\!\sim} [\,|\, w_r\colon \mathit{secretary}\ \sigma_6,\ w_r\colon o_5\ \mathit{owns}\ \sigma_6\,]$

f. $B \oplus \beta_{21} \oplus \beta_{16} \oplus [\,|\, w_r\colon o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,] \mathrel{\models\!\!\!\sim}$
$[\,|\, \mathbf{not}(\beta_{25} \oplus [\,|\, \mathbf{not}[\,u_{12} \,|\, w_r\colon o_8\ \mathit{have}\, u_{12},\ w_r\colon \mathit{secretary}\ u_{12}\,])\,]$

g. $B \oplus \beta_{21} \oplus \beta_{16} \not\models\!\!\!\sim [\,|\, w_r\colon o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,]$

h. $B \oplus \beta_{21} \oplus \beta_{16} \not\models\!\!\!\sim [\,|\, \mathbf{not}[\,|\, w_r\colon o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,]\,]$

It may be observed that the use of the name 'Mark' leads to the introduction of a fresh discourse marker with the second occurrence. With the first use of a name in a discourse, the participants (or arbitrary hearers) select a referent directly in the state of affairs the discourse is about. The referent must be an individual which their world knowledge tells them is known to all of them, is known to carry this name, and which is the most plausibe candidate given what is stated about him or her in the discourse. Known individuals are easily uniquely identifiable given the situations and events in which they figure. This carries over to the second use of a name in the discourse, so that on the force of the common knowledge of the participants, and the speaker's duty to provide sufficient information for the hearer to be able to identify the individuals he refers to uniquely, it is likely that they will pick the same referent with repeated use of the name.[57]

This does not fully explain, however, why the same referent is picked unfailingly with the repeated use of the name in discourse. There is something of a discourse effect going on, which ensures local continuity of reference. We suggest

[56] For some motivation of this requirement, see Van Leusen (2004).

[57] This also explains why there are relatively few misunderstandings between the participants in a conversation about who is the intended referent of a given name.

that names introduce topical individuals, and that they come with the additional requirement that they should uniquely identify an individual in their local topic domain. In the case of example (6.87) the discourse relation *denial* may be taken to set the topic domain of its right argument to the topic domain of the left argument, so that the denial and the statement denied have a common topic. The uniqueness requirement on the name 'Mark' relative to the local topic domain then enforces that o_8 and o_9 have the same occupant in the state of evaluation. The interpreting agent may accept $[\,|\, w_r : o_8 \textbf{ is } o_9\,] \sqsubseteq B$.

What about the presuppositional condition (6.91e) triggered by 'Mark's secretary'? The denial results in a discourse meaning which entails that Mark does not have a secretary. The only way to satisfy both Global Consistency and the presuppositional condition is by accommodating the background material in β_{16}: $[\,\sigma_6 \,|\, w_r\!: \mathit{secretary}\ \sigma_6,\ w_r\!: o_5\ \mathit{owns}\ \sigma_6\,] \sqsubseteq \beta_{16}$. Setting β_{16} to $[\,\sigma_6 \,|\, w_r\!: \mathit{secretary}\ \sigma_6,\ w_r\!: o_5\ \mathit{owns}\ \sigma_6\,]$ and β_{21} and β_{25} to $[\,|\,]$ on the basis of the preference for smallest possible background values, the interpreting agent may infer discourse meaning (6.92a), where (6.92b-e) are satisfied. (6.92c) is the incompatibility condition triggered by the *denial* relation. Clearly, this is satisfied. The local informativity and consistency conditions on the left argument of the denial (6.92d) and (6.92e) can be satisfied as before.

(6.92)a. $B \oplus [\,|\, \textbf{not}[\,\sigma_6 \,|\, w_r\!: \mathit{secretary}\ \sigma_6,\ w_r\!: o_5\ \mathit{owns}\ \sigma_6,\ w_r\!: o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,],\ \textbf{not}[\,u_{12} \,|\, w_r\!: o_8\ \mathit{have}\, u_{12},\ w_r\!: \mathit{secretary}\ u_{12}\,]\,]$

b. $[w_r\, o_0\, o_5\, o_8 \mid w_r\!: \mathit{John}\ o_0,\ w_r\!: \mathit{Mark}\ o_5,\ w_r\!: \mathit{Mark}\ o_8,\ w_r : o_8 \textbf{ is } o_9] \sqsubseteq B$

c. $B \oplus [\,\sigma_6 \,|\, w_r\!: \mathit{secretary}\ \sigma_6,\ w_r\!: o_5\ \mathit{owns}\ \sigma_6,\ w_r\!: o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,] \models [\,|\, \textbf{not}[\,|\, \textbf{not}[\,u_{12} \,|\, w_r\!: o_8\ \mathit{have}\, u_{12},\ w_r\!: \mathit{secretary}\ u_{12}\,]\,]\,]$

d. $B \oplus [\,\sigma_6 \,|\, w_r\!: \mathit{secretary}\ \sigma_6,\ w_r\!: o_5\ \mathit{owns}\ \sigma_6\,] \not\models [\,|\, w_r\!: o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,]$

e. $B \oplus [\,\sigma_6 \,|\, w_r\!: \mathit{secretary}\ \sigma_6,\ w_r\!: o_5\ \mathit{owns}\ \sigma_6\,] \not\models [\,|\, \textbf{not}[\,|\, w_r\!: o_0\ \mathit{made\text{-}app\text{-}with}\ \sigma_6\,]\,]$

In short, in interpreting B's reply participant A finds out that global accommodation of the presupposition is no longer an option: the information that Mark has a secretary contradicts the semantic content of B's contribution, resulting in an inconsistent discourse meaning. Global Consistency prohibits this. A reading is selected so that the background material is captured locally, in the scope of the negation introduced by the denial. In the course of the exchange, the hearer's model of the discourse meaning changes from one that entails that Mark has a secretary to one that entails that Mark does not have a secretary. In this sense the presupposition evoked by the first speaker's utterance is *cancelled* through the denial.

As with the disjunctive discourse connective, a justification condition might be assigned to the right argument of the denial operator. In the example at hand, where β_{25} is set to $[\,|\,]$, it would be trivially satisfied. The justification axiom $\mathcal{A}27$

triggers a nonsatisfaction constraint on the local background value of the left argument, β_{16}. This is instantiated as $B \not\models [\,\sigma_6 \,|\, w_r\!:\mathit{secretary}\ \sigma_6,\ w_r\!:o_5\ \mathit{owns}\ \sigma_6\,]$ which given Global Consistency and the discourse meaning (6.92a)must be satisfied. Axiom $\mathcal{A}27$ enforces that background material accommodated in the scope of the denial is noninferrable outside its scope. Conversely, presuppositional conditions introduced through the statement that is denied which may plausibly be satisfied outside the scope of the denial will *not* be accommodated within its scope.

6.6 Conclusion

The aim we set ourselves was to integrate two approaches to presupposition theory in our framework of discourse interpretation. One we dubbed the 'satisfaction' approach and the other the 'anaphoric approach' to presupposition interpretation. The central claim of the first is that a presupposition evoked by a given clause must be entailed or satisfied in the local context of the clause. The initial context of interpretation may be strengthened in order to satisfy a presupposition; this is global accommodation. The main point of the other approach is that presuppositions are anaphoric; if they can be bound in the embedding context, the descriptive material carried by the presupposition trigger must be inserted where the binder sits, otherwise, the presupposition is accommodated. This means that the descriptive material (including a binder) is inserted in an appropriate place in the discourse meaning.

On the assumption that these approaches capture different, but equally important aspects of presupposition theory we provided a combined approach in which, on the one hand, presupposition triggers can be analysed as anaphors, and on the other, they introduce a context-sensitive condition which says that the descriptive material they carry must be entailed in their local context. Presupposition interpretation is modeled as a reasoning process in which binding and satisfaction go hand in hand. If a trigger is anaphoric, the presuppositional condition must be satisfied given the binding of the anaphoric marker or markers in the presuppositional material. Conversely, the presuppositional condition functions as a suitability condition on the resolution of the anaphoric marker. Although we concentrated on the analysis of definite descriptions, an immediate advantage of our combined treatment is that it allows a unified analysis of various types of presupposition triggers, while individual differences can be accounted for.

Fusing the two approaches, we found that each comes with its own type of accommodation. Starting out from the perspective of satisfaction theory, we defined accommodation as the specification of general background knowledge in the implicit background of the conversation B. We called this global accommodation. In order to satisfy presuppositional conditions, world knowledge may freely be accepted in B—in as far as this is compatible with the other constraints and pref-

erences in the grammar, of course. Different from what we find in the anaphoric approach, what is accommodated need not be precisely the presuppositional material. It can be any information which the interpreting agent is willing to accept as common knowledge, either permanently or temporarily. What is accepted in B will count as information shared between himself and the other discussion participants.

While global accommodation accounts for strengthening of the initial context of interpretation, it cannot account for cases in which the interpretation process calls for the strengthening of local contexts other than B and results in the modification of the semantic content of a discourse. Most authors agree that there is a class of 'nonpresupposing' readings for presupposition evoking sentences which can best be accounted for if we accept that binders and presuppositional material can be accommodated in subordinated positions in a discourse meaning. In the anaphoric approach this is easily accounted for. There, accommodation is the insertion of presuppositional material in a location on the accessibility path upward from the presupposition trigger in the DRS that represents the semantic content of a discourse.

We incorporated the anaphoric notion of accommodation in our grammar by generalising the original treatment over *local background values*, associating these with all discourse units in a discourse description. To avoid unwarranted weakening of discourse meanings, local backgrounds, though underspecified, are required to be composed of linguistically generated material only. By restricting the background material to presuppositional material and binders, we obtain a constructive treatment of accommodation as in the anaphoric approach to presupposition. This form of accommodation was called content modification. The outcome, in our combined approach, is that we have two qualitatively different forms of accommodation. Both are the effect of the hearer's effort to obtain a possible and most preferred discourse meaning; if the interpretation task calls for it, they may co-occur.

Presupposition theory, as implemented in our framework of discourse interpretation, is an integrated part of discourse theory. When a sentence contains a presupposition trigger, obviously the satisfaction of the presuppositional condition is a factor in its interpretation. It is not the only factor by far, however. The hearer may use the total of his grammatical knowledge combined with what the discourse description tells him, to infer the possible grammatical readings of a sentence. Reasoning on the basis of his preferences and world knowledge, he will determine the most preferred readings of the sentence. What comes out as a most preferred reading is accounted for in terms of the interaction of all grammatical constraints and preferences involved. This explains that sentences that may be said to evoke weak presuppositions can get strongly presupposing interpretations.

An important restriction on what may count as acceptable or felicitous readings is represented by the conditions on consistency and informativity in our grammar. We distinguished global consistency and informativity conditions from

local ones. The first constrain output discourse meanings, the second constrain the semantic values of discourse units relative to their local context. The acceptability conditions on presupposition resolution of the anaphoric approach can be modeled in LDG by means of these conditions. It was shown that, while on the one hand, local informativity and consistency requirements account for various cases of nonpresupposing discourse, on the other, associating them uniformly with all discourse units results in a discourse theory that is too strong.

To remedy this, it was suggested that local informativity and consistency conditions are 'soft constraints', which can be violated without resulting in completely infelicitous discourse. This places local informativity and consistency conditions in the realm of pragmatics in our discourse grammar. Readings which violate the conditions are no longer predicted to be ungrammatical. Rather, they are predicted to be dispreferred, and may come in various degrees of infelicity in principle. Soft constraints on individual discourse units would follow from a general preference for more informative discourse. It remains to be seen how the testing of soft constrains can be implemented technically in our grammar.

Preferences play a prominent role in the account in general. They are essential in cutting down the number of possible readings of a given sentence. Both the satisfaction approach and the anaphoric approach to presupposition make reference to the interaction with preferences and world knowledge. Though we have not provided a formal treatment of the preference system, we may claim that 'reasoning with preferences' finds a place naturally in our framework of discourse interpretation. It supplements logical deduction and common sense reasoning, and interacts freely with the interpreting agent's world knowledge, and syntactic and semantic knowledge in the interpretation task.

We adopted a general preference for more plausible readings, a preference for smaller background values (economy), and a preference for accommodation in the main DRS. The last one corresponds to the preference for global over nonglobal accommodation in the anaphoric account. The preference for 'binding' over 'accommodation' in the anaphoric account follows from the second one. Various candidate preferences popped up in the discussion of the examples: a preference for logically stronger readings was briefly considered, a preference for avoiding ambiguity was suggested, and a preference for selecting readings reflecting the speaker's point of view (contra hearer's world knowledge) was mentioned. It was hinted that when interpreting discourse in an out-of-the-blue context, hearers preferably set the parameters for 'genre' or 'conversation type' to unspecific, default values.

As for the predictions resulting from our combined treatment of presupposition, if presupposition triggers are analysed as anaphors, and if the preferences in the discourse grammar are set in the manner of the anaphoric account, the predictions conform to those of the anaphoric approach. The explanatory force of the satisfaction approach is preserved in the assumption that presupposition triggers introduce context sensitive conditions: presuppositions are *tests* on the

local context of the expression that triggers them. As such, they set going the reasoning process that eventually results in the binding of presuppositional material, either by an explicit antecedent, or by accommodated material.

Content modification, however, takes away part of the explanatory power of the presuppositional conditions: under certain conditions, presuppositions are satisfied by their own presuppositional material. As became clear in the course of our investigations, this can be remedied by adding what we have called 'justification conditions' on background material. Justification conditions explain why background material is inserted where it is. In as far as we have specified them, they come partly as dependency constraints associated with individual discourse operators, and partly in the form of an informativity constraint on the local context of the incorporating operator.

If we accept justification conditions in our discourse theory, the overall picture that emerges from our treatment is one in which, one the one hand, presupposition interpretation is constructive to the extent that it results in the insertion of presuppositional material at a location in the discourse tree representation. On the other, the selection of a most suitable location for the projected material is controlled by satisfaction and non-satisfaction conditions on background material that sits in these locations. We leave as a question for future research whether justification conditions could possibly replace the presuppositional condition, given some leaner set-up of our treatment where the background material carried by a trigger corresponds exactly to the presuppositional material.

Justification conditions, as they were implemented in our grammar, do not replace the informativity and consistency constraints on semantic values mentioned earlier. They express that background information must be rooted locally, in its context of interpretation. This is most clearly illustrated with bridging anaphors, where justification conditions induce the computation of the bridging link. In their way, justification conditions contribute to the establishment of cohesion, and the selection of cohesive readings of discourse as the most preferred ones. The same goes for the presuppositional conditions associated with triggers. In fact, the two factors that are combined in our treatment of presupposition, namely anaphoricity and inference in local contexts, represent main sources of cohesion in discourse.

We conclude that the satisfaction approach to presupposition and the anaphoric approach can fruitfully be combined in a single framework of discourse interpretation, if the discourse semantics is rich and fine-grained enough to describe constraints on binding as well as conditions on local contexts. The approaches are not incompatible, and with the right setting of preferences and constraints in the grammar, can be seen to strengthen each other. Integrating them in a full-fledged discourse theory reduces the work load of presupposition theory as such. Anaphoric aspects of presuppositions can be explained in terms of the interaction with independent conditions on the properness of discourse meanings, and topicality and salience requirements on their binders. The analysis of pre-

suppositions as conditions on local contexts fits in with the general treatment of context-sensitivity in our grammar framework. The interaction with preferences and the hearer's nonlinguistic knowledge in the interpretation task are part and parcel of the framework.

The treatment of presupposition that we provided can be worked out further, and as we tried to indicate, in various spots alternative solutions can be chosen. We have been more concerned with showing some different corners of the search space, than with putting forward the one and only correct presupposition theory. We hope that on the basis of a careful consideration of that search space, motivated choices can eventually be made. This may concern, for instance, the treatment of so-called 'intermediate readings' occurring with conditionals and quantificational structures. Or it may involve making a choice in the precise lexical description of definite descriptions, with a view to the treatment of bridging anaphors. It may consist in the further spelling out and testing of justification conditions on background material, and possibly, providing some refinement on the preference for smaller background values that we adopted. An account of the costliness of accommodation may be provided.

The coverage of the theory could be extended by the investigation of other presupposition triggers then the definite determiner, both with respect to their anaphoricity and their accommodation potential. Furthermore, it might be investigated whether the local backgrounds, that are now used only to capture material introduced by presupposition triggers, can be employed in, for instance, focus interpretation, and the computation of parallellism. Finally, some general parameters of the discourse theory must be specified. A notorious candidate for this, by now, is discourse availability, or the theory of discourse topicality and salience. Another is the preference system.

Chapter 7

Conclusion

In this work, the LDG framework of discourse interpretation was presented. The framework treats linguistic analysis as a part of human reasoning. It combines a Logical Description Grammar for discourse with a reservoir of world and common sense knowledge, and a preference system. While the grammar represents a language user's grammatical or linguistic knowledge, the reservoir of world and common sense knowledge stands for his nonlinguistic knowledge. The preference system represents his capacity to draw default inferences on the basis of his linguistic and nonlinguistic knowledge, and weigh possible readings of discourse relative to each other, in order to determine most preferred ones.

The grammar formalism, developed in Van Leusen and Muskens (2003) building on Vijay-Shanker (1992), Muskens (1995; 2001), and Gardent and Webber (1998) presupposes two levels of linguistic analysis: the level of discourse descriptions, and the level of the tree structures or linguistic representations that verify these descriptions. The processing of a discourse T results in a *discourse description* Δ_T, which is a set of statements in a logical language representing the language user's knowledge of the discourse. The logical language or *description language* is type logic. The grammar $\mathcal{G}$ is a set of statements in type logic as well. The combined logical theory $\mathcal{G} + \Delta_T$ describes the syntax and semantics of T in parallel. A language user obtains verifying trees by reasoning about Δ_T given $\mathcal{G}$. Thus, *parsing* is logical deduction. Descriptions can be partial. Since $\mathcal{G} + \Delta_T$ can have more than one model, the syntax and semantics of T may remain underspecified. Discourse is processed incrementally, and the incrementation operation is a monotonic, that is, no information is ever retracted from a discourse description.

Linguistic representations of discourse are tree structures, decorated with syntactic and semantic information. The semantic information at each node consists in a semantic value, a local context composed in the style of Karttunen (1974), and a background value. The semantic representation language is fine-grained CDRT, a variant of compositional DRT which is spelled out in Van Leusen and Muskens (2003). The semantic value of the root is constructed compositionally

from the semantic values and background values of the nodes it dominates. The *discourse meaning* is defined as the update of what is taken for granted in the implicit general background of the discourse with the semantic value of the root node.

Sentence grammar is fully integrated with discourse grammar. As in Lexicalised Tree Adjoining Grammar the lexicon is the 'motor' of the grammar system. All the nodes, structure, and semantic values occurring in a discourse tree are contributed by lexemes or implicit discourse relations described in the lexicon, and occurring in the discourse. Beside a unified treatment of sentence and discourse analysis, the LDG architecture guarantees a smooth integration of grammatical levels. The grammar describes syntax and semantics in tamdem, and the preference system covers pragmatics aspects of interpretation. The language user may access all knowledge sources in parallel.

As a neutral, constraint based system, the framework is a useful tool for the specification of linguistic theory. While, obviously, we have made certain choices in the set-up of the discourse theory, individual parameters can be set differently, even where basic matters like the syntactic structure of discourse representations or the semantic representation language are concerned. The framework can thus be employed to integrate and compare different approaches towards one and the same topic in discourse theory. To show the flexibility of the discourse model, it was applied in the domain of anaphoricity and presupposition theory.

We found that with a discourse semantics that is sufficiently fine-grained and rich, approaches to presupposition projection as different in perspective as the so-called *satisfaction approach* to presupposition (Karttunen 1974; Heim 1983; Beaver 1995) and the *anaphoric approach* (Van der Sandt 1992; Geurts 1999) can be integrated in a single treatment. The first treats presuppositions as tests on the context of interpretation. For a sentence to be felicitous or interpretable, the context of interpretation must entail its presuppositions. Sometimes a presupposition can only be satisfied by strengthening the context of interpretation; this is accommodation in the sense of the satisfaction approach. The central claim of the second approach is that presuppositional material must be bound in the discourse context of the trigger. If it cannot be bound, it must be inserted at some semantically accessible location in the discourse context; this is accommodation in the sense of the anaphoric approach.

We put forward a treatment of anaphoricity in which anaphoric elements introduce underspecified discourse markers. Through properness conditions on local contexts and the discourse meaning, the grammar enforces that these discourse markers are bound. Anaphora resolution is a reasoning process, in which a language user determines possible binders for a given anaphoric marker and selects solutions that result in a most preferred reading of the discourse. We propose that presuppositional elements in the language introduce a condition that requires the *entailment* of the descriptive material they carry in their local context. The relations of entailment and inclusion are made available at object level

in the semantics. By letting presupposition triggers introduce presuppositional conditions while at the same time analysing them as anaphors, the gist of both approaches to presupposition projection is captured.

Presupposition theory is part of discourse theory. The two approaches are formulated in terms of the interaction of general parameters in the discourse theory. Two parameters that obviously play a role are *anaphoricity*, and *givenness* or inferrability of information in the context of interpretation. A third factor in the account is *accommodation*, that is, the specification of background information. A final one is the constraining effect of *preferences*. Main preferences in our treatment concern the consistency and informativity of discourse units relative to their local context, the plausibility of resulting readings, a preference for accommodation at top level in the discourse meaning, and 'economy' or the measure to which readings require accommodation at all.

Accommodation comes in two forms in our treatment. One is called 'global accommodation', and consists in the specification of the implicit general background of the discourse; in line with the satisfaction approach common knowledge may freely be accepted in the implicit background if the interpretation task calls for it. The other is called 'content modification', and consists in the insertion of background material in locations in the discourse meaning that are accessible from the presupposition trigger. In line with the anaphoric approach what is accommodated is precisely the presuppositional material, including binders. Our treatment predicts that global accommodation and content modification may co-occur.

The predictions of our combined presupposition theory are variable, because they depend on the setting of preferences, and choices in the implementation of the discourse theory in general. If we analyse presupposition triggers as anaphors, the predictions conform to those of the anaphoric approach, though generally speaking we can distinguish more readings. This is due to the fact that our treatment of accommodation is richer, and more of the account is handled with preferences. What comes out as a most preferred reading is the result of the interaction of all grammatical constraints and preferences involved. Thus, sentences that may be said to evoke weak presuppositions can get strongly presupposing interpretations.

We conclude that the satisfaction approach and the anaphoric approach to presupposition are not incompatible. In terms of predictions they overlap and in terms of explanatory force they can be seen to strengthen each other. This is particularly clear with accommodation, where 'satisfaction type' global accommodation may co-occur with 'anaphoric style' content modification. Some spin-off of our treatment in an underspecification formalism is that a neat declarative implementation could be provided of aspects of presupposition projection which are often handled procedurally in the anaphoric approach to presupposition. Furthermore, though our focus was on definite descriptions, because of the flexibility and generality of our analysis we expect that it can easily be extended to other classes of presupposition triggers. Next, due to the fact that *inferrability* in the local context can be stated as a characterising condition of context-sensitive ele-

ments in the language a straightforward treatment of bridging anaphors could be provided. Finally, we gained some insight in the role of informativity and consistency constraints in discourse theory. It was tentatively concluded that these can better be handled by means of the preference system or soft constraints than by means of hard constraints.

Much remains to be done, both where the specification of discourse theory and where the formal properties of the grammar system are concerned. Our treatment of presupposition theory can be further investigated and consolidated. For instance, the role of so-called 'justification conditions', which govern the insertion of background material in accessible locations, could be further explored. The computation of parallelism and a treatment of focus interpretation might be spelled out, and its interaction with global accommodation and content modification in our grammar investigated. A theory of topic-comment structure and its effect on the topic hierarchy of discourse needs to be specified. In terms of this the parameter of discourse availability, topicality, and salience could be filled in, and its interaction with various classes of anaphors made concrete.

The grammar can be rolled out in different directions. For one, a component of phonology and intonation needs to be specified. Furthermore, the semantics can be extended by providing an event semantic fragment and treatment of intensionality in fine-grained CDRT. A semantics for interrogative sentences might be integrated in the grammar, and we could try to deal with the semantics and pragmatics of questions and answers. The semantics of argumentative discourse and nonmonotonic discourse contributions could be further explored. With this topic, which essentially concerns the commitments and beliefs of the participants, we probably cross the limit between what can be treated at object level in discourse descriptions, and what must be treated at meta level. Generally speaking, we expect that the specification of a 'grammar of discourse contributions' or a treatment of speech act theory will be a meta level affair.

Technically, there are various questions waiting to be answered. First, we have seen that preferences play a crucial role in our explanations. We are in the process of uncovering what preferences are relevant, and of delimiting aspects in discourse theory that can be handled by hard constraints from those that require a treatment with soft constraints. We assumed that a language user can reason with preferences just as with his other knowledge sources but we have not provided an implementation of this. A treatment along the lines of Optimality Theory[1] could be pursued. Furthermore, computional aspects of the formalism, in particular of reasoning with discourse descriptions need to be explored. Conceptually, the grammar models the interaction between different grammatical levels transparantly, but can reasoning about syntax and semantics in tandem in type logic also be implemented efficiently? What will be the effect on the computa-

[1]See Prince and Smolensky (1997), Gilbers and de Hoop (1998), Blutner and Zeevat (2004), and Blutner et al. (2006) on Optimality Theory.

tion of verifying models for discourse descriptions if we allow less strict linguistic structures than tree structures?[2] And what is the effect of choosing a different semantic representation language? In general, how can the intertwining of linguistic parsing and logical inference best be tackled in practical implementations?

Furthermore, formal aspects of the grammar as an incremental system may be explored. In our discourse model, language users reason about the verifying trees of discourse descriptions given temporary right-closure of the description. It is clear that when the right-closure assumption is dropped, part of the inferences or conclusions drawn from the description are lost. One would like to know what part of the conclusions drawn is independent of the right-closure assumption and are preserved from update to update. Of the other inferences one would like to be able to say which will be preserved by default across incrementation steps.

I would like to end here with a plea for synthesis and generality in discourse theory. Much of the current work in descriptive sentence and discourse semantics is highly specialised in terms of the issues it addresses, and there is a great variety in the formal representation systems employed. Still it is tacitly agreed upon that all findings in a given natural language are to be explained from a single model of a language user's linguistic capacity. An isolated analysis of, say, the word 'again' in English, can hardly be evaluated on its merits if it is not an integrated part of discourse theory, where we can see how it interacts with presupposition theory, focus interpretation, and tense, to name some relevant parameters. Analyses of different topics in linguistics must be combined so that they may help to evaluate and strengthen each other. But even when such interactions are taken into account, our analyses remain gratuitous if there is no common base of assumptions that we all agree upon. I think that common base is present, but it certainly takes an effort to dig it up from the diversity of formal systems and competing theories. The description grammar method may bring some unity in diversity, and help to specify our discourse theories in one and the same language.

[2]Compare Blackburn and Gardent (1998) and Duchier and Gardent (1999) on the computation of minimal models when the verifying structures are DAGs.

Appendix A

A Fine-grained Compositional DRT

In the main text a formalism for semantic representation was introduced in a somewhat loose fashion. The purpose of this appendix is to give a tighter formalisation of this system.

Recall that the *description language* that was chosen for the grammar is classical type logic $\mathcal{L}$, and that $\mathcal{L}$ is used to talk about all aspects of linguistic structure including semantic values. Hence, it figures as the semantic representation language as well. Our task thus is not to introduce a new logic, but to specifically tune $\mathcal{L}$ to the description of semantic values and semantic relations. In fact, $\mathcal{L}$ (plus some axioms) will be used to talk about semantic mechanisms that are normally only described metalinguistically. In Muskens (2001) it was shown how some essential mechanisms of predicate logic can be mimicked on the level of the object language, using a technique virtually identical to the technique used for obtaining the embedding of DRT in $\mathcal{L}$ known as compositional discourse representation theory or CDRT (Muskens 1996). In this appendix, which is based upon Van Leusen and Muskens (2003), we introduce a treatment of DRT in type logic which is similar to CDRT but offers a somewhat more fine-grained notion of DRSs. We shall sometimes refer to it as 'fine-grained CDRT' or CDRT$^+$.

Typographical Conventions Let us start with an overview of the typographical conventions that will be used. We work in classical type logic with ground types e (entities), ν (tree nodes), l (tree labels), π (registers or discourse markers), and s (states). It will be expedient to have *pairing* and *projection* in the logic. We will assume that whenever A is a term of type α and B a term of type β, $\langle A, B\rangle$ is a term of type $\alpha \times \beta$. Also, whenever A is a term of type $\alpha \times \beta$, $\mathsf{fst}(A)$ will be a term of type α (denoting the first element of the denotation of A) and $\mathsf{snd}(A)$ a term of type β (denoting the second element). These terms will have the obvious semantics.

Since we will have occasion to use many variables and constants in the ground types that were mentioned and in complex types composed out of these primitive

Variable	Type	Constants	Type	Constants	Type
x, y, z	e	r, 0, 1, ...	ν	$\prec$, $\lhd^+$	$(\nu \times \nu)t$
k	ν	PEDRO, HAS, ...	νt	V	$(\pi \times s)t$
v	π	*elaboration, ...*	νt	E	$(e \times e)t$
i, j	s	*lex, imp, term*	νt	W	et
c	$st \times \pi t$	*pn, cn, tv, rel, ...*	νt	*Pedro, ...*	$(e \times e)t$
		dp, vp, d, ...	l	*sleeps, ...*	$(e \times e)t$
		ℓ	νl	*loves, has, ...*	$(e \times e \times e)t$
		α^+, α^-	$\nu\nu$	B	τ
		u, o, w	$\nu\pi$	Γ	$\nu\tau$
		σ^β	$\nu\beta$		

Table A.1: Some variables and constants with their associated types. Here β varies over types and τ is an abbreviation of $\pi t \times (st \times \pi t)t$ (the type of DRSs).

ones, it will be convenient to have conventions of the form: 'whenever v is used, it will be a variable of type π.' Table A.1 lists these and also gives the types of certain constants that will be employed. (Why these types are chosen will hopefully become clear below.)

Binding Since the treatment of variables (discourse referents) and binding in Discourse Representation Theory, the theory we wish to emulate internally, is fundamentally different from their treatment in classical logic, it will be necessary to axiomatize it here. This is done with axioms $\mathcal{A}28$, $\mathcal{A}29$ and $\mathcal{A}30$ below (see Muskens 1991, 1996, 2001 for more information[1]). Variables will be modeled as *registers*, memory locations that can hold certain values. They will have type π. A *state* (type s) intuitively is an assignment of values (type e) to each register. We write $V(v, i)$ for 'the value of register v in state i' and if $\delta_1, \ldots, \delta_n$ are terms of type π, we write $i[\delta_1 \ldots \delta_n]j$ for

$$\forall v \left[(v \neq \delta_1 \wedge \ldots \wedge v \neq \delta_n) \rightarrow V(v', i) = V(v', j)\right] ,$$

i.e. for 'i and j differ at most in $\delta_1, \ldots, \delta_n$' (with v the first variable in some fixed ordering which is not free in $\delta_1, \ldots, \delta_n$). The first axiom, $\mathcal{A}28$, says that, in each state, each register can be updated selectively, i.e. its value can be set to any x while the values of other registers remain unchanged. This axiom makes states and registers essentially behave as assignments and variables respectively.

$\mathcal{A}28$ $\forall i \forall v \forall x \, \exists j \, [i[v]j \wedge V(v, j) = x]$

$\mathcal{A}29$ $\forall k k' \, \rho(k) \neq \rho'(k')$, if $\rho, \rho' \in \{u, o, w\}$ are distinct

[1]Note that the axiom numbers here differ from the ones employed in these papers

$\mathcal{A}30$ $\forall k_1 k_2 \, [\rho(k_1) = \rho(k_2) \rightarrow k_1 = k_2]$, if $\rho \in \{u, o, w\}$

We let u, o, and w be functions from nodes to registers. Axiom $\mathcal{A}29$ ensures that the images of these functions are disjoint, while $\mathcal{A}30$ requires each to be injective. In this way we make sure that fresh registers come with each node. This in turn will make it easy to let indefinites be associated with new discourse referents. The values of u will typically be associated with *new* referents, while the values of o are referents that belong to the 'background' of the discourse. The values of w are registers that can store *worlds* (see below). In this dissertation we will only use $w(r)$, the w value of the root.[2] A notational convention that we find useful is to write arguments of u, o, and w as subscripts, e.g. u_3 instead of $u(3)$, w_r instead of $w(r)$ etc.

Worlds While states correspond to assignments and registers correspond to variables, the notion of a possible world corresponds to the technical notion of a *model*. We will make some use of possible worlds here, but will not change the logic—a step often associated with their introduction. Instead we will consider possible worlds to be objects of type e, i.e. we will simply take them to be (abstract) entities. The predicate letter W will be used for the predicate 'is a world' and E will be an existence predicate, so that $E(x, y)$ stands for 'object x exists in world y' and $\lambda x \; E(x, y)$ is y's domain. In our semantics we shall make use of a (finite) set $\mathcal{V}$ of predicate letters ($\{$*Pedro, has, mule, feeds, ...*$\}$) all of whose arguments are of type e and whose last arguments are to be interpreted as worlds. For instance, $has(x, y, z)$ should be read as 'x has y in world z'. The following axiom scheme requires that last arguments indeed are worlds and, somewhat rigorously, demands that all other arguments of an $\mathcal{V}$ relation denote objects in this world's domain.[3]

$\mathcal{A}31$ $\forall x_1 \ldots x_n \, y \, [R(x_1, \ldots, x_n, y) \rightarrow (W(y) \wedge E(x_1, y) \wedge \ldots \wedge E(x_n, y))]$,
for each $R \in \mathcal{V}$

A Principle of Plenitude While $\mathcal{A}28$ in a sense required there to be enough states present for the binding mechanism to work, the axiom scheme $\mathcal{A}32$ below will put similar requirements on worlds. We will ensure that there are enough worlds, our models should contain at least all worlds that correspond to finite models for the language $\mathcal{V}$.[4] Define a $\mathcal{V}$-*atom* to be an atomic formula $R(t_1, \ldots, t_n)$

[2]But in extensions containing modal operators it would be natural to use other values of w as well.

[3]This will serve present purposes. That the requirement needs to be fine-tuned can be seen from such classic examples such as 'worship', where the object argument need not exist. The axiom also will need to be adapted if accessibility relations between worlds are considered, something we shall not do here.

[4]Compare Van Benthem (1986), where it is argued that, in order to keep complexity down, natural language semantics should concentrate on finite models.

with $R \in \mathcal{V}$. A $\mathcal{V}$-*literal* is a $\mathcal{V}$-atom or the negation of a $\mathcal{V}$-atom. The axiom scheme may be instantiated by any number of variables $x_1, \ldots, x_n$ and any conjunction of $\mathcal{V}$-literals $\varphi(x_1, \ldots, x_n, y)$ satisfying the conditions stated.

$\mathcal{A}32$ $\exists x_1 \ldots x_n y\,[W(y) \wedge x_1 \neq x_2 \wedge x_1 \neq x_3 \wedge \ldots \wedge x_{n-1} \neq x_n \wedge$
$\forall x\,[E(x,y) \leftrightarrow (x = x_1 \vee \ldots \vee x = x_n)] \wedge \varphi(x_1, \ldots, x_n, y)]$,
where $\varphi(x_1, \ldots, x_n, y)$ is any conjunction of $\mathcal{V}$-literals with at most $x_1, \ldots, x_n$ and y free that does not contain both a $\mathcal{V}$-atom and its negation.

This axiom scheme will have instantiations such as, say,

$$\begin{aligned}&\exists x_1 x_2 x_3 y[W(y) \wedge x_1 \neq x_2 \wedge x_1 \neq x_3 \wedge x_2 \neq x_3\\ &\wedge \forall x\,[E(x,y) \leftrightarrow (x = x_1 \vee x = x_2 \vee x = x_3)]\\ &\wedge \mathit{has}(x_1, x_2, y) \wedge \neg \mathit{has}(x_2, x_3, y) \wedge \mathit{mule}(x_2, y) \wedge \mathit{hay}(x_3, y)]\ .\end{aligned}$$

Thus every finite world is stipulated to exist, as we can easily produce an exhaustive description in this way. Note that this does not only ensure the existence of many worlds but also the existence of a multiplicity of states. Since, according to $\mathcal{A}28$, each register of a state can be selectively updated with type e values, and since worlds are type e objects, each state will have many variants differing only from it in the value of the register w_r.[5]

A Fine-grained Compositional DRT In the following paragraphs we will give some definitions that will make key concepts of Discourse Representation Theory available in the description language of the discourse grammar. The definitions are based upon the semantic axioms that were given above and the DRT concepts that result are used throughout the book on what might be called the semantic tier of the grammar formalism.

Compositional Discourse Representation Theory (Muskens 1996; 2001) is an emulation of DRT in type logic. Semantic representations in CDRT look like DRSs, but actually they are abbreviations of formulae in type logic mimicking the behaviour of DRT. An essential characteristic of the CDRT formalism is that $\mathcal{L}$ is used to talk about notions which in standard DRT belong to the metalanguage. We have already seen how assignments (verifying embeddings) and variable binding can be modeled at object level and will see some more examples shortly. Some of the structure present in the *representations* of standard DRT is will be imported to the level of the semantic objects they denote.

The approach chosen here has much in common with that of Muskens (1996), but also differs from it in an important respect. What the two approaches have in common is that both are based on a transcription of truth conditions for DRT,

[5]States that have worlds in u or o registers are not formally excluded, but are of no relevance. We will also not bother to formally exclude that worlds are elements of the domains of (other) worlds.

carried out within an axiomatic extension of type logic. The difference between the two treatments results from the fact that while in Muskens (1996) the DRT truth conditions that were transcribed were those of Groenendijk and Stokhof (1991), we use a variant of the semantics in Zeevat (1989) here. In this way we get a more fine-grained DRT semantics (in the sense that we get stronger requirements on the identity of DRSs), which suits our present purposes better and may be thought interesting in itself.

Preliminary Syntactic Sugaring Some abbreviations introducing set-theoretic notation will come in handy.

A.1. Definition. Let $A_1, \ldots, A_n$ be terms of some type α and let X be the first variable of that type not free in these terms, then

$$\{A_1, \ldots, A_n\} \quad \text{abbreviates} \quad \lambda X\,[X = A_1 \vee \ldots \vee X = A_n]\ .$$

Furthermore, let A and B be terms of some type $(\alpha_1(\cdots(\alpha_n t)\cdots))$ and let $X_1, \ldots, X_n$ be variables such that each X_i is of tye α_i. Then

$$\begin{aligned} A \cup B \quad &\text{abbreviates} \quad \lambda X_1 \ldots X_n\,[AX_1 \ldots X_n \vee BX_1 \ldots X_n]\ , \\ A \cap B \quad &\text{abbreviates} \quad \lambda X_1 \ldots X_n\,[AX_1 \ldots X_n \wedge BX_1 \ldots X_n]\ , \\ A \subseteq B \quad &\text{abbreviates} \quad \forall X_1 \ldots X_n\,[AX_1 \ldots X_n \rightarrow BX_1 \ldots X_n]\ . \end{aligned}$$

Choosing Types for Conditions and DRSs In Zeevat (1989) the semantic value of a Discourse Representation Structure K is a pair consisting of (a) a set of discourse referents (the universe of K) and (b) a set of assignments. The latter consists of all those assignments that verify all conditions in K. This could easily be transposed to the present set-up by letting the first part of any DRS K consist of a set of registers (type πt) and its second part of a set of states (type st), so that the type of DRSs would be $\pi t \times st$. However, we opt for a variant of this approach and will deviate from Zeevat's treatment in two respects. Firstly, we will follow Visser (1994) in keeping track of the *free* discourse referents in any DRS or condition (see also Van Eijck and Kamp (1997)). This can be done by letting conditions be pairs of (a) sets of states and (b) sets of registers (those that are free in the condition). The type of a condition will then become $st \times \pi t$.

The second deviation from Zeevat is that we will let the second part of a DRS K consist of a *set of* conditions rather than of a single condition representing their conjunction. Having sets of conditions around without conjoining their members embodies the hypothesis that the language system sometimes addresses one of these in isolation.

With these modifications the type of DRSs becomes $\pi t \times (st \times \pi t)t$, which we will often abbreviate as τ. The alternative formalization will give stronger identity criteria on DRSs.

Freedom and Properness Given that the free referents of a condition form its second element, the free referents of any DRS K can easily be computed: these are the ones that are free in some condition of K but are not in the universe of K. A DRS will be *proper* if it has no free referents. These notions are made available to the logic with the following definitions.

$$\begin{array}{rll} \mathit{free}(K) & \text{abbreviates} & \lambda v \exists c\,[\neg \mathsf{fst}(K)(v) \wedge \mathsf{snd}(K)(c) \wedge \mathsf{snd}(c)(v)] \\ \mathit{proper}(K) & \text{abbreviates} & \neg \exists v\, \mathit{free}(K)(v) \end{array}$$

Boolean Structure on DRSs The two following notions, *merge* and *inclusion*, play a key role in the algebra of Discourse Representation Structures.

A.2. Definition. Let K and K' be terms of type τ.

$$\begin{array}{rll} K \oplus K' & \text{abbreviates} & \langle \mathsf{fst}(K) \cup \mathsf{fst}(K'), \mathsf{snd}(K) \cup \mathsf{snd}(K') \rangle \\ K \sqsubseteq K' & \text{abbreviates} & \mathsf{fst}(K) \subseteq \mathsf{fst}(K') \wedge \mathsf{snd}(K) \subseteq \mathsf{snd}(K') \end{array}$$

We will typically be interested in type τ objects $\langle \{\delta_1, \ldots, \delta_n\}, \{\gamma_1, \ldots, \gamma_m\} \rangle$ with a finite universe $\{\delta_1, \ldots, \delta_n\}$ and a finite set of conditions $\{\gamma_1, \ldots, \gamma_m\}$. Such DRSs we prefer to write as $[\delta_1 \ldots \delta_n \mid \gamma_1, \ldots, \gamma_m]$, in a way reminiscent of the usual 'box' notation for DRSs.

Truth and Consequence A DRS K is *true* in a state i if there is a j, differing from i at most as far as the universe of K is concerned, such that j verifies all conditions in K. We will write $\mathit{true}(K)$ for

$$\lambda i \exists j (\forall v\,[\neg \mathsf{fst}(K)(v) \rightarrow V(v,i) = V(v,j)] \wedge \forall c(\mathsf{snd}(K)(c) \rightarrow \mathsf{fst}(c)(j))) \ .$$

Here $\forall v\,[\neg \mathsf{fst}(K)(v) \rightarrow V(v,i) = V(v,j)]$ says that i and j differ at most with respect to discourse referents in the universe of K, while $\forall c(\mathsf{snd}(K)(c) \rightarrow \mathsf{fst}(c)(j)))$ says that j satisfies (the first parts of) all conditions of K.

On the finite DRSs we are interested in, $\mathit{true}(K)$ is equivalent with the more readable

$$\lambda i \exists j\ (i[\delta_1 \ldots \delta_n] j \wedge \mathsf{fst}(\gamma_1)(j) \wedge \ldots \wedge \mathsf{fst}(\gamma_m)(j)) \ .$$

The notion of DRS truth immediately leads to a notion of DRS consequence: K' is said to *follow from* K if $\mathcal{A} \models \mathit{true}(K) \subseteq \mathit{true}(K')$, i.e. if K' is true in all states i in which K is true, in any model of the axioms. This notion can be relativised to any description $\mathcal{D}$ by saying that K' *follows from* K *given* $\mathcal{D}$ if $\mathcal{D}, \mathcal{A} \models \mathit{true}(K) \subseteq \mathit{true}(K')$. It will be seen shortly how a similar notion can be obtained on the level of the describing logic.

The DRT Sublanguage Until now we have not paid much attention to the internal structure of DRSs, but this will now be remedied by giving a definition of the DRT language as a sublanguage of type theory. The next definition gives notation for conditions.

A.3. DEFINITION. Let $R \in \mathcal{V}$, let $\delta_1, \ldots, \delta_n$ be terms of type π (discourse referents), and let K and K' be terms of type τ (DRSs). Then

$$\begin{array}{rll} R\{\delta_1, \ldots, \delta_n\} & \text{abbreviates} & \langle \lambda i.R(V(\delta_1, i), \ldots, V(\delta_n, i)), \{\delta_1, \ldots, \delta_n\}\rangle \\ \delta_1 \textbf{ is } \delta_2 & \text{abbreviates} & \langle \lambda i.V(\delta_1, i) = V(\delta_2, i), \{\delta_1, \delta_2\}\rangle \\ \textbf{not } K & \text{abbreviates} & \langle \lambda i.\neg \mathit{true}(K)(i), \mathit{free}(K)\rangle \\ K \Rightarrow K' & \text{abbreviates} & \textbf{not } (K \oplus [\ |\ \textbf{not } K']) \\ K \textbf{ or } K' & \text{abbreviates} & \langle \mathit{true}(K) \cup \mathit{true}(K'), \mathit{free}(K) \cup \mathit{free}(K')\rangle \end{array}$$

Note that $R\{\delta_1, \ldots, \delta_n\}$ predicates R of the *values* of $\delta_1, \ldots, \delta_n$ in some state, not of these registers themselves. We shall write $R\{\delta, w_r\}$ as $w_r\colon R\delta$ and $R\{\delta_1, \delta_2, w_r\}$ as $w_r\colon \delta_1 R \delta_2$. It is also worthwile to note that the definition keeps track of the referents that are free in any condition.

Let us consider the sublanguage of type $st \times \pi t$ and type τ terms that is given by the following Backus-Naur form, and dub it the *DRT sublanguage*. Here the δ range over π terms that have the form $\rho(n)$, where $\rho \in \{u, o, w\}$ and n is a node name. The R are taken from $\mathcal{V}$.

$$\begin{array}{rcl} \gamma & ::= & R\{\delta_1, \ldots, \delta_n\} \mid \delta_1 \textbf{ is } \delta_2 \mid \textbf{not } K \mid K \Rightarrow K' \mid K \textbf{ or } K' \\ K & ::= & [\delta_1 \ldots \delta_n \mid \gamma_1, \ldots, \gamma_m] \end{array}$$

Translation into Predicate Logic It is useful to know that there are simple truth-preserving translations from the DRT sublanguage to predicate logic. The following one is taken from Muskens (1996) (but see also Kamp and Reyle (1993)).

A.4. DEFINITION. Let $\dagger$ be a function that injectively maps each δ to a variable of type e. The function **tr** from $st \times \pi t$ terms in the DRT sublanguage to predicate logical formulas and the function **wp**, which takes a pair consisting of a τ term and a predicate logical formula and yields a predicate logical formula, are defined as follows.

$$\begin{array}{rcl} \mathbf{tr}(R\{\delta_1, \ldots, \delta_n\}) & = & R(\delta_1^\dagger, \ldots, \delta_n^\dagger) \\ \mathbf{tr}(\delta_1 \textbf{ is } \delta_2) & = & \delta_1^\dagger = \delta_2^\dagger \\ \mathbf{tr}(\textbf{not } K) & = & \neg \mathbf{wp}(K, \top) \\ \mathbf{tr}(K \Rightarrow K') & = & \neg \mathbf{wp}(K, \neg \mathbf{wp}(K', \top)) \\ \mathbf{tr}(K \textbf{ or } K') & = & \mathbf{wp}(K, \top) \vee \mathbf{wp}(K', \top) \\ \mathbf{wp}([\delta_1 \ldots \delta_n \mid \gamma_1, \ldots, \gamma_m], \chi) & = & \exists \delta_1^\dagger \ldots \delta_n^\dagger [\mathbf{tr}(\gamma_1) \wedge \ldots \wedge \mathbf{tr}(\gamma_m) \wedge \chi] \end{array}$$

For any formula φ and state variable i, let φ^i be the result of substituting $V(\delta, i)$ for $\delta^\dagger$, for each $\delta^\dagger$ that is free in φ. The following holds.

A.5. THEOREM. *Let K be a τ term in the DRT sublanguage and let i be an arbitrary type s variable. Let Σ contain all statements $\delta \neq \delta'$, for every pair δ, δ' of syntactically different discourse referents in K. Then*

$$\Sigma, \mathcal{A} \models \mathbf{wp}(K, \top)^i \leftrightarrow \mathit{true}(K)(i)$$

The disequality statements are needed: If, say, 12 and 9 corefer, it follows that $[u_{12} \mid w_r\colon \mathit{mule}\ u_9]$ and $[u_9 \mid w_r\colon \mathit{mule}\ u_9]$ corefer as well, but the translation would not preserve this. In practice this will be no limitation at all.

For a proof of the theorem, abbreviate $\lambda i \lambda j\, (i[\delta_1 \ldots \delta_n]j \wedge \mathsf{fst}(\gamma_1)(j) \wedge \ldots \wedge \mathsf{fst}(\gamma_m)(j))$ as $gs(K)(i)(j)$ and use induction on complexity in the DRT sublanguage construction to show that, given the axioms and disequalities, $\mathbf{wp}(K, \chi)^i \leftrightarrow \exists j\, [gs(K)(i)(j) \wedge \chi^j]$ and $\mathbf{tr}(\gamma)^i \leftrightarrow \gamma(i)$, for all K and γ. This can be shown using the methods of Muskens (1996, 2001). $gs(K)(i)(j)$ is the *Groenendijk-Stokhof* notion of DRS update potential.

Note that, given these rules, $\mathbf{wp}([\ \mid\ ], \top) = \top$, i.e. the empty DRS $[\ \mid\ ]$ will be true in every state. Conversely $\mathbf{wp}([\ \mid \mathbf{not}[\ \mid\ ]], \top) = \bot$, so that $[\ \mid \mathbf{not}[\ \mid\ ]]$ in fact plays the role of the falsum on the DRS level. We will sometimes abuse notation and write $\top$ for $[\ \mid\ ]$ and $\bot$ for $[\ \mid \mathbf{not}[\ \mid\ ]]$.

Object Level Entailment Logical relations such as entailment and consistency can be approximated *at the level of the describing logic*, as we shall see presently. Suppose that K is a DRS with w_r in its universe, while K' does not have w_r in its universe. Then the condition $K \Rightarrow K'$ quantifies over possible worlds. This is best illustrated with an example: let us take K to be $[w_r\, o_1 \mid w_r\colon \mathit{mule}\ o_1]$, 'there is a mule', and let K' be $[u_2 \mid w_r\colon \mathit{farmer}\ u_2,\ w_r\colon u_2\ \mathit{owns}\ o_1]$, 'some farmer owns the mule', so that K' should not be made to follow from K. Then, for arbitrary i,[6]

$$\mathbf{tr}(K \Rightarrow K')^i = \forall y \forall x_1\, [\mathit{mule}(x_1, y) \rightarrow \exists x_2\, [\mathit{farmer}(x_2, y) \wedge \mathit{owns}(x_2, x_1, y)]]\ .$$

In view of $\mathcal{A}31$ this is equivalent with

$$\begin{array}{l} \forall y[W(y) \rightarrow \\ \quad \forall x_1\, [(\mathit{mule}(x_1, y) \wedge E(x_1, y)) \rightarrow \\ \quad\quad \exists x_2\, [E(x_2, y) \wedge \mathit{farmer}(x_2, y) \wedge \mathit{owns}(x_2, x_1, y)]]]\ , \end{array}$$

from which it is seen that indeed quantification over worlds is involved. Note that $\mathcal{A}32$ immediately provides a singleton-domain counterexample to this purported

[6] Here $K \Rightarrow K'$ is closed. This will be the typical situation in applications.

strict implication and in general it will give finite counterexamples if there are such.

Conditions $K \Rightarrow K'$ are of type $st \times \pi t$ and we will typically want to state conditions on the t level. The following definition adds the necessary quantification over states / assignments, but also excludes that the condition $K \Rightarrow K'$ contains free discourse referents.

A.6. Definition. Let K and K' be terms of type τ.

$$K \mathrel{|\!\sim} K' \quad \text{abbreviates} \quad \forall i\ \mathsf{fst}(K \Rightarrow K')(i) \wedge \neg\exists v\ \mathsf{snd}(K \Rightarrow K')(v)$$

We will write $K \not\mathrel{|\!\sim} K'$ for $\neg(K \mathrel{|\!\sim} K')$. Note that $\mathrel{|\!\sim}$ in a sense only approximates the notion of ordinary entailment: not all possible counterexample worlds are guaranteed to exist in any model, although all finite ones are.

Object Level Substitution It will come in handy to have available a substitution operation on discourse referents. We define $\{\delta/\delta'\}K$, 'the *substitution* of δ' by δ in a DRS K' as follows.

A.7. Definition. Let U range over terms of type πt (sets of discourse referents), γ over terms of type $st \times \pi t$ (conditions), C over terms of type $(st \times \pi t)t$ (sets of conditions), and K over terms of type $\pi t \times (st \times \pi t)t$ (DRSs). We write

$$\begin{array}{rll}
\{\delta/\delta'\}U & \text{for} & \lambda v[(U(v) \wedge \neg v = \delta') \vee v = \delta] \\
\{\delta/\delta'\}\gamma & \text{for} & \langle \lambda i \exists j[i[\delta']j \wedge V(j,\delta') = V(i,\delta) \wedge \mathsf{fst}(\gamma)(j)], \{\delta/\delta'\}\mathsf{snd}(\gamma)\rangle \\
\{\delta/\delta'\}C & \text{for} & \lambda c \exists c'[C(c') \wedge c = \{\delta/\delta'\}c'] \\
\{\delta/\delta'\}K & \text{for} & \langle \{\delta/\delta'\}\mathsf{fst}(K), \{\delta/\delta'\}\mathsf{snd}(K)\rangle
\end{array}$$

Given this definition, for instance $\{u_2/u_1\}[\,u_1 \mid w_r\colon \mathit{sleep}\, u_1\,]$ is equivalent to $[\,u_2 \mid w_r\colon \mathit{sleep}\, u_2\,]$.

The Background As explained in the main text, the background B should have at least three properties. First, it should be *proper*, secondly, none of the discourse referents u_k, which should be *novel*, should already be contained in its universe, but, thirdly, the actual world w_r should be so contained. The following axiom ensures this.

$\mathcal{A}19$ $\mathit{proper}(B) \wedge \neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$

Appendix B

Axioms of the Grammar

This appendix sums up the general descriptions or axioms in the $\mathcal{A}$ component of our LDG grammar systems, apart from the ones that define the CDRT$^+$ sub-language of our description language, which can be found in appendix A. We also include the main definitions of incremental LDG here.

Dominance and precedence in tree structures In the description language of the grammar nodes are indicated by $k, k_1, k_2, \ldots$. Furthermore $\prec$ denotes *linear precedence*, $\lhd$ denotes *immediate dominance*, $\lhd^*$ denotes *dominance*, and $\lhd^+$ *proper dominance*. A node k_1 dominates a node k_2 if k_1 properly dominates k_2, or if k_1 is identical to k_2.

$\mathcal{A}1$ $\forall k\, r \lhd^* k$

$\mathcal{A}2$ $\forall k \neg k \lhd^+ k$

$\mathcal{A}3$ $\forall k_1 k_2 k_3\, [[k_1 \lhd^+ k_2 \wedge k_2 \lhd^+ k_3] \rightarrow k_1 \lhd^+ k_3]$

$\mathcal{A}4$ $\forall k \neg k \prec k$

$\mathcal{A}5$ $\forall k_1 k_2 k_3\, [[k_1 \prec k_2 \wedge k_2 \prec k_3] \rightarrow k_1 \prec k_3]$

$\mathcal{A}6$ $\forall k_1 k_2\, [k_1 \prec k_2 \vee k_2 \prec k_1 \vee k_1 \lhd^+ k_2 \vee k_2 \lhd^+ k_1 \vee k_1 = k_2]$

$\mathcal{A}7$ $\forall k_1 k_2 k_3\, [[k_1 \lhd^+ k_2 \wedge k_1 \prec k_3] \rightarrow k_2 \prec k_3]$

$\mathcal{A}8$ $\forall k_1 k_2 k_3\, [[k_1 \lhd^+ k_2 \wedge k_3 \prec k_1] \rightarrow k_3 \prec k_2]$

$\mathcal{A}9$ $\forall k_1 k_2\, [k_1 \lhd k_2 \rightarrow k_1 \lhd^+ k_2]$

$\mathcal{A}10$ $\forall k_1 k_2 k_3 \neg [k_1 \lhd k_3 \wedge k_1 \lhd^+ k_2 \wedge k_2 \lhd^+ k_3]$

Axiom $\mathcal{A}1$ states that the root r dominates all other nodes. $\mathcal{A}2$–$\mathcal{A}5$ ensure that $\lhd^+$ and $\prec$ are strict partial orderings. The *Exhaustivity* property, $\mathcal{A}6$, says that any two nodes stand either in a precedence relation or in a relation of dominance.The *Inheritance* property, laid down in $\mathcal{A}7$ and $\mathcal{A}8$, says that precedence relations are inherited from ancestor nodes. Finally, $\mathcal{A}9$ and $\mathcal{A}10$ constrain the relation between immediate dominance and proper dominance.

Syntactic Labelling and Lexical Anchoring The description language contains the *labeling function* ℓ, which assigns syntactic category labels to nodes. To ensure that nodes are labeled uniquely, it must be prevented that category label names corefer. This is what $\mathcal{A}11$ does.

$\mathcal{A}11$ $c_1 \neq c_2$, if c_1 and c_2 are distinct label names

The following three general descriptions are the *anchoring* axioms of the grammar. The predicate *lex* says that a node is lexical, and the functions α^+ and α^- produce, respectively, the positive and negative anchor of a node. For example $\alpha^+(k_1) = k_2$ says that k_1 is positively anchored to k_2.

$\mathcal{A}12$ $\forall k\, lex(\alpha^+(k))$

$\mathcal{A}13$ $\forall k\, [k = r \vee lex(\alpha^-(k))]$

$\mathcal{A}14$ $\alpha^-(r) = r$

Axiom $\mathcal{A}12$ requires that the positive anchor of a node is lexical and $\mathcal{A}13$ requires the same for negative anchoring, except for the root. The root is negatively anchored to itself ($\mathcal{A}14$). The positive and the negative anchor of a given node can but need not be identical.

Incorporating Implicit Anchors To incorporate implicit anchors we differentiate between lexically expressed nodes and terminal nodes. *Terminal nodes* are nodes which do not dominate any other node:

$\mathcal{A}15$ $\forall k[term(k) \leftrightarrow \neg\exists k' : k \lhd^+ k']$

Anchors are now required to be terminal, rather than lexical. We adjust the anchoring axioms $\mathcal{A}12$ and $\mathcal{A}13$, replacing '*lex*' with '*term*'.

$\mathcal{A}12'$ $\forall k : term(\alpha^+(k))$

$\mathcal{A}13'$ $\forall k[k = r \vee term(\alpha^-(k))]$

The predicate *imp* indicates that a node is terminal but not lexical. This makes the node into an *implicit anchor*:

$\mathcal{A}16$ $\forall k[imp(k) \leftrightarrow term(k) \wedge \neg lex(k)]$

Implicit anchors occur only with discourse relations:

$\mathcal{A}17$ $\forall k : imp(k) \rightarrow rel(k)$

Incremental LDG The relation of *immediate precedence*, denoted by $\prec_l$ and the relation of *immediate lexical succession (ILS)*, denoted by the symbol $\prec_\iota$ are defined as follows.

Immediate Precedence $\forall k_1, k_2[k_1 \prec_l k_2 \leftrightarrow k_1 \prec k_2 \wedge \neg \exists k_3(k_1 \prec k_3 \prec k_2)]$

ILS $\forall k_1, k_2[k_1 \prec_\iota k_2 \leftrightarrow (k_1 \prec_l k_2 \vee \exists k_3(k_1 \prec_l k_3 \prec_l k_2 \wedge imp(k_3)))]$

A node k_2 *immediately lexically succeeds* a node k_1, or k_2 is a *immediate lexical successor* of k_1 if and only if k_1 immediately precedes k_2, or there is a single implicit anchor which immediately succeeds k_1 and immediately precedes k_2.

The characterising properties of the incrementation operation are *Addition* and *Lexical Contiguity*. The first adds the new input description δ_{i+1} to the discourse description built up so far, Δ_i. The latter makes the leftmost lexical node named in the new input description into the immediate lexical successor of the rightmost lexical node named in the discourse description. The initial discourse description is the empty set.

Addition $\Delta_i \cup \{\delta_{i+1}\} \subseteq \Delta_i + \delta_{i+1}$.

Lexical Contiguity $\{n_{RML\Delta_i} \prec_\iota n_{LML\delta_{i+1}}\} \subset \Delta_i + \delta_{i+1}$, where $i > 0$.

Initial Discourse Description $\Delta_0 = \emptyset$

Incrementation

1. $\Delta_0 + \delta_1 = \Delta_0 \cup \{\delta_1\}$
2. $\Delta_i + \delta_{i+1} = \Delta_i \cup \{\delta_{i+1}\} \cup \{k_{RML\Delta_i} \prec_\iota k_{LML\delta_{i+1}}\}$, where $i > 0$.

Left and Right Closure To compute verifying tree models of discourse descriptions, they must temporarily be closed off. The following axioms constrain 'left-closure' (LC) and 'right-closure' (RC) of discourse descriptions. The special node names 0 and n_{fin} refer to the nodes representing the first and the last lexical element in a discourse. n_{fin} changes with each incrementation step.

LC $\neg \exists k(k \prec 0)$

RC $\neg \exists k(n_{fin} \prec k)$

Local Contexts and the Implicit Background For each node k in a tree structure, its *local context* $\Gamma(k)$ is computed top-down from the local context of its mother node (if it has one) and the semantic values of its (left) sister nodes. This is regulated by means of the elementary structures in the lexicon and axiom $\mathcal{A}18$. The axiom sets the local context of the root node $\Gamma(r)$ to the implicit background or global context of the discourse B (a type τ constant), and it says that nodes of a syntactic category other than S inherit the local context of their mother node. Here $k_1 \lhd k_2$ abbreviates $k_1 \lhd^+ k_2 \wedge \neg\exists k[k_1 \lhd^+ k \wedge k \lhd^+ k_2]$.

$\mathcal{A}18\ \ \Gamma(r) = B\ \wedge\ \forall k_1 k_2\,[[k_1 \lhd k_2 \wedge \ell(k_2) \neq s] \rightarrow \Gamma(k_2) = \Gamma(k_1)]$

The background B is subject to a few general constraints, and these are summed up in $\mathcal{A}19$. We use $\mathsf{fst}(B)$ to denote the universe of B. The implicit background must be a proper DRS, its universe should not contain any u-markers, and the marker occupied by the actual world, w_r, must be bound in the universe of B.

$\mathcal{A}19\ \ \mathit{proper}(B) \wedge \neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$

No node can have w_r as its semantics:

$\mathcal{A}20\ \ \neg\exists k\ \sigma_k^\pi = w_r$

Felicity Conditions on Consistent and Informative Updating These come in two forms: global conditions and local ones. The first constrain discourse meaning and may be seen as characterising conditions of discourse contributions. The second apply to discourse constituents. They constrain the arguments of individual discourse connectives and relations.

Global Consistency $B \oplus \sigma_r \not\models \bot$

Global Informativity For a discourse contribution updating a discourse description Δ to $\Delta + \delta$, there must be a DRS K such that

$\Delta^{rcl} \not\models (B \oplus \sigma_r \mid\!\sim K)$ and $(\Delta + \delta)^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$.

Local Informativity $\forall k\,[\,\ell(k) = s \rightarrow \Gamma_k \not\models \sigma_k]$

Local Consistency $\forall k\,[\,\ell(k) = s \rightarrow \Gamma_k \not\models [\,|\ \mathbf{not}\ \sigma_k]\,]$

Discourse contributions may be classified in terms of the monotonicity property:

Monotonic contribution A discourse contribution resulting in an update $\Delta + \delta$ of a discourse description Δ is *monotonic* iff
$\forall K$ if $\Delta^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$ then $(\Delta + \delta)^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$

Nonmonotonic contribution A discourse contribution resulting in an update $\Delta + \delta$ of a discourse description Δ is *nonmonotonic* iff
$\exists K$ such that $\Delta^{rcl} \models (B \oplus \sigma_r \mid\!\sim K)$ and $(\Delta + \delta)^{rcl} \not\models (B \oplus \sigma_r \mid\!\sim K)$

Properness Condition For any discourse description, the discourse meaning $B \oplus \sigma_k$, and the local contexts of all nodes in the described structure must be proper DRSs:

$\mathcal{A}21$ $\mathit{proper}(B \oplus \sigma_r) \wedge \forall k[\, \mathit{proper}(\Gamma_k)]$

Axiom $\mathcal{A}21$ subsumes the first part of axiom $\mathcal{A}19$, which states some properties of the implicit background. The latter may now be reformulated as

$\mathcal{A}19'$ $\neg\exists k\, \mathsf{fst}(B)(u_k) \wedge \mathsf{fst}(B)(w_r)$

Chomskian Binding Conditions Using $\mathsf{gvc}(k)$ to represent the governing category of a node k, and $\mathsf{cc}(k, k')$ to denote that a node k c-commands a node k', the binding principles (A), (B), and (C) can be formalized as additional constraints within the discourse theory, as follows.

Binding Conditions in LDG

(A) $\forall k[\, \mathit{refl}(k) \rightarrow \exists k'(\mathsf{gvc}(k) \lhd^* k' \wedge \mathsf{cc}(k', k) \ \wedge\ \sigma(k) = \sigma(k'))\,]$

(B) $\forall k[\, \mathit{pron}(k) \rightarrow \neg\exists k'(\mathsf{gvc}(k) \lhd^* k' \wedge \mathsf{cc}(k', k) \wedge \sigma(k) = \sigma(k'))\,]$

(C) $\forall k[\, \mathit{pn}(k) \vee \mathit{det}(k) \rightarrow \neg\exists k'(\mathsf{cc}(k', k) \wedge \sigma(k) = \sigma(k'))\,]$

Referential and Existential Readings An alternative treatment may be adopted in which some discourse referents declared in the universe of the global background are referential while others are interpreted existentially. This can be achieved by employing the existing difference between o-markers and u-markers to distinguish between referential and existential interpretations, rather than 'background' and 'new' markers. The following semantic axiom regulates this; '*ref*' is a πt-type predicate which marks a discourse referent as referential.

$\mathcal{A}22$ $\forall k.\mathit{ref}(o_k) \wedge \forall k.\neg \mathit{ref}(u_k)$

We no longer exclude u-markers from the universe of the global background of the conversation. Arguably, it may also be required that o-markers, resulting in referential interpretations, are in the universe of the global background only. Thus, instead of $\mathcal{A}19$ we would have

$\mathcal{A}19''$ $\mathit{proper}(B) \wedge \forall k\, \mathsf{fst}(B)(o_k) \wedge \mathsf{fst}(B)(w_r)$

In addition, the truth definition we had must be replaced with one that excludes the set of all o-markers from interpretation as an existentially bound variable: 'a DRS K is said to be *true in state i with respect to a set of externally anchored referents* '*ref*' if $\mathit{true}(\langle \mathsf{fst}(K) - \mathit{ref},\ \mathsf{snd}(K)\rangle)(i)$ holds.' As before, o-markers receive their interpretation from i_0 directly.

Content Modification Every S-node k in a discourse description comes with a *background set* $\mathbf{b}(k)$ of type νt, a possibly empty set of nodes. The nodes function as pointers to linguistically generated material. In principle, any lexical element or implicit anchor with accommodation potential can provide this material, but we focus on anaphors and presupposition triggers here. These come with a value Π_k of type τ which constitutes the *background material* generated by the trigger. For every node k, a *local background* or *background value* β_k of type τ is constructed from its background set $\mathbf{b}(k)$ by taking the generalised merge of the background material at the nodes in $\mathbf{b}(k)$:[1]

$\mathcal{A}23\ \forall k[\,\beta_k = \bigoplus\{\Pi_{k'} \mid k' \in \mathbf{b}(k)\}\,]$

In order to account for intermediate accommodation in conditionals and quantifier structures, it may be assumed that S-nodes come, in addition, with a subsidiary background set $\bar{\mathbf{b}}(k)$ of type νt. Axiom $\mathcal{A}19$ must then be strengthened to $\mathcal{A}23'$.

$\mathcal{A}23'\ \forall k[\,\beta_k = \bigoplus\{\Pi_{k'} \mid k' \in \mathbf{b}(k)\} \wedge \bar{\beta_k} = \bigoplus\{\Pi_{k'} \mid k' \in \bar{\mathbf{b}}(k)\}\,]$

Local backgrounds must be composed of presuppositional material generated by the stretch of discourse they are the local background of. The following axiom takes care of this. It says that background sets $\mathbf{b}(k)$, or $\bar{\mathbf{b}}(k)$, contain only terminal nodes k' dominated by k. The effect of this is that the background value of an arbitrary node or discourse unit k will be constructed only from background material introduced by elements in the lexical yield of k.

$\mathcal{A}24\ \forall k, k'[k' \in \mathbf{b}(k) \vee k' \in \bar{\mathbf{b}}(k)) \rightarrow (term(k)\ \wedge\ k \lhd^* k')]$

Arguably, background material, if it is accommodated, can only be consumed once. To avoid multiple accommodation, axiom $\mathcal{A}25$ may be adopted.

$\mathcal{A}25\ \forall k, k', k''[(k' \in \mathbf{b}(k) \wedge k' \in \mathbf{b}(k'')) \rightarrow k'' = k]$ and
$\forall k, k', k''[(k' \in \bar{\mathbf{b}}(k) \leftrightarrow k' \notin \mathbf{b}(k)) \wedge ((k' \in \bar{\mathbf{b}}(k) \wedge k' \in \bar{\mathbf{b}}(k'')) \rightarrow k'' = k)]$

For each S-node k, its *information content* i_k is the merge of its background value and its semantic value:

$\mathcal{A}26\ k[\ell(k) = s\ \rightarrow\ i_k = \beta_k \oplus \sigma_k]$

Discourse meanings may now be described as the update of the general background B with the information content of the discourse: $B \oplus i_r$.

[1] This abbreviates the statement
$\forall k[\forall v(\mathsf{fst}(\beta(k))(v) \leftrightarrow \exists k'(\mathbf{b}(k)(k') \wedge \mathsf{fst}(\Pi(k'))(v))) \wedge$
$\forall c(\mathsf{snd}(\beta(k))(c) \leftrightarrow \exists k'(\mathbf{b}(k)(k') \wedge \mathsf{snd}(\Pi(k'))(c)))]$

For any given S-node k, its local context is constructed top down from the local context of its mother node, the information content of left sister nodes when present, and its own local background. Most of this is laid down in the lexicon of the grammar. The following axiom, which is a revised version of $\mathcal{A}18$, sets the local context of the root node r to $B \oplus \beta_r$ and the local contexts of non S-nodes to the local context of their mother node.

$$\mathcal{A}18'\ \ \Gamma(r) = B \oplus \beta_r \ \wedge\ \forall k_1 k_2\,[[k_1 \lhd k_2 \wedge \ell(k_2) \neq s] \rightarrow \Gamma(k_2) = \Gamma(k_1)]$$

Economy The following preference constrains the local background values of nodes. The effect is that readings resulting from nontrivial content modification are predicted to be less preferred, or 'more costly', than those that do not require content modification at all. The Van der Sandtian preference for 'binding' over 'accommodation' follows from this.

[**Prefer smaller local backgrounds**]
For all nodes k and DRSs K, K', prefer $\beta_k = K$ over $\beta_k = K'$ if $K \sqsubseteq K'$.

Justification Condition The following axiom requires that background material is accepted in the scope of an operator only when it is *informative* outside the scope of the operator. We generally refer to the combination of constraints triggered by $\mathcal{A}27$ plus the dependency constraints on background values associated with individual operators in the lexicon as *justification* conditions. They justify the insertion of background material in a certain location, when content modification takes place.

$$\mathcal{A}27\ \ \forall k[\,(k \neq r \wedge \beta(k) \neq [\,|\,])\ \rightarrow\ \Gamma(\mathit{top}(\alpha^-(k)) \not\models \beta(k)\,]$$

We use $\mathit{top}(k)$ to indicate the topnode in the elementary structure anchored by node k:

$$\forall k, k'[\,\mathit{top}(k) = k' \leftrightarrow (k = \alpha^+(k') \wedge \neg\exists k''[k = \alpha^+(k'') \wedge k'' \lhd^+ k'])\,]$$

Bibliography

Abeillé, A. (1988a). A lexicalised tree adjoining grammar for French: The general framework. Technical report, University of Pennsylvania.

Abeillé, A. (1988b). Parsing French with tree adjoining grammar: some linguistic accounts. In *Proceedings of COLING*, Budapest, pp. 7–12.

Abeillé, A. (1993). *Les Nouvelles Syntaxes. Grammaires d'unification et analyse du français*. Paris: Armand Colin.

Abeillé, A. (1994). Syntax or semantics? handling non-local dependencies with mctags or synchronous tags. *Computational Intelligence 4*(10), 471–485.

Ariel, M. (1990). *Accessing Noun-Phrase Antecedents*. London: Routledge.

Asher, N. (1993). *Reference to Abstract Objects in Discourse*, Volume 50 of *Studies in Linguistics and Philosophy*. Dordrecht: Kluwer Academic Publishers.

Asher, N. (1995). From discourse macro-structure to micro-structure and back again: Discourse semantics and the focus/background distinction. In H. Kamp and B. Partee (Eds.), *Proceedings of the Conference on Semantics in Context*, SFB340 Report. University of Stuttgart.

Asher, N. and A. Lascarides (1998a). Bridging. *Journal of Semantics 15*(1), 83–113.

Asher, N. and A. Lascarides (1998b). The semantics and pragmatics of presupposition. *Journal of Semantics 3*(15), 239–300.

Asher, N. and A. Lascarides (2003). *Logics of Conversation*. Cambridge: Cambridge University Press.

Austin, J. (1961). Truth. In J. Urmson and G. Warnock (Eds.), *Philosophical Papers*, pp. 117–133. Oxford: Oxford University Press.

Austin, J. (1962). *How to Do Things with Words*. Oxford: Clarendon Press.

Barwise, J. and J. Perry (1983). *Situations and Attitudes*. Cambridge, Massachusetts: MIT Press.

Beaver, D. (1994). Accommodating topics. ILLC Prepublication Series LP-94-17, University of Amsterdam.

Beaver, D. (1995). *Presupposition and Assertion in Dynamic Semantics*. Ph. D. thesis, University of Edinburgh.

Beaver, D. (1996). Local satisfaction preferred. In P. Dekker and M. Stokhof (Eds.), *Proceedings of the Tenth Amsterdam Colloquium*. ILLC/ University of Amsterdam.

Beaver, D. I. (1997). Presupposition. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, Chapter 17, pp. 939–1008. Elsevier Science B.V.

Beaver, D. I. (2001). *Presupposition and Assertion in Dynamic Semantics*. Stanford, California: CSLI Publications.

van Benthem, J. (1986). *Essays in Logical Semantics*. Reidel.

Blackburn, P. and J. Bos (1999). *Representation and Inference for Natural Language: An Advanced Course in Computational Semantics*, Volume II, Working with Discourse Representation Structures. CSLI Publications. Draft.

Blackburn, P. and J. Bos (2005). *Representation and Inference for Natural Language: A First Course in Computational Semantics*, Volume I. CSLI Publications.

Blackburn, P. and C. Gardent (1998). A specification language for discourse. In *Proceedings of LACL'98 (Logical Aspects of Computational Linguistics*, Grenoble, France. Internet version is called: a description language for discourse semantics.

Blutner, R. (2000). Some Aspects of Optimality in Natural Language. *Journal of Semantics 17*(3), 189–216.

Blutner, R., H. de Hoop, and P. Hendriks (2006). *Optimal Communication*. Stanford: CSLI Publications.

Blutner, R. and H. Zeevat (Eds.) (2004). *Optimality Theory and Pragmatics*. Houndmills, Basingstoke, Hampshire: Palgrave/Macmillan.

Bos, J. (2003). Implementing the Binding and Accommodation Theory for Anaphora Resolution and Presupposition Projection. *Computational Linguistics 29*(2), 179–210.

Bos, J. (2004). Computational semantics in discourse: Underspecification, resolution, and inference. *Journal of Logic, Language and Information 13*, 139 –175.

Carlson, L. (1983). *Dialogue Games*. Dordrecht: Reidel.

Charniak, E. (1983). Passing markers: a theory of contextual influence in language comprehension. *Cognitive Science 7*, 171–190.

Chen, J. and K. Vijay-Shankar (1995). Towards a reduced commitment, d-theory style tag parser. In *Proceedings of the International Workshop on Parsing Technologies*.

Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht: Foris.

Cimiano, P. and U. Reyle (2005). Talking about Trees, Scope and Concepts. In H. Bunt, J. Geertzen, and E. Thijsse (Eds.), *Proceedings of IWCS-6*, Tilburg, the Netherlands, pp. 90–102.

Clark, H. (1977). Bridging. In P. N. Johnson-Laird and P. C. Wason (Eds.), *Thinking: Readings in Cognitive Science*. Cambridge University Press.

Cooper, R. (1979). The Interpretation of Pronouns. In F. Heny and H. Schnelle (Eds.), *Selections from the Third Groningen Round Table: Syntax and Semantics*, Volume 10, pp. 61–92. Academic Press.

Crain, S. and M. Steedman (1985). On not being led up the garden-path: The use of context by the psychological parser. In D. Dowty, L. Kartunnen, and A. Zwicky (Eds.), *Natural Language Parsing: Psychological, Computational and Theoretical Perspectives*. Cambridge University Press.

Dahl, O. (1974). Topic-comment structure revisited. In *Topic and Comment, Contextual Boundness and Focus*. Hamburg: Buske.

Dalrymple, M., M. Kanazawa, Y. Kim, S. Mchombo, and S. Peters (1998). Reciprocal expressions and the concept of reciprocracy. *Linguistics and Philosophy 21*, 159–210.

Dalrymple, M., M. Kanazawa, S. Mchombo, and S. Peters (1994). What do reciprocals mean? In M. Harvey and L. Santelmann (Eds.), *Proceedings of Semantics and Linguistic Theory (SALT) IV*. Cornell University.

van Deemter, K. (1989). Anaphora across the Board. Draft. To appear in Linguistics and Philosophy.

van Deemter, K. (1990). Forward References in Natural Language. *Journal of Semantics 7*, 281–300.

van Deemter, K. and E. Krahmer (1998). On the interpretation of anaphoric noun phrases: Towards a full understanding of partial matches. *Journal of Semantics 15*(2), 355–392.

Dekker, P. (1993). *Transsentential Meditations: Ups and Downs in Dynamic Semantics*. Ph. D. thesis, Universiteit van Amsterdam.

Donnellan, K. (1966). Reference and Definite Descriptions. *Philosophical Review 75*, 281–304. (Reprinted from Philosophical Review).

Duchier, D. and C. Gardent (1999). A constraint-based treatment of descriptions. In *Third International Workshop on Computational Semantics IWCS-3*, pp. 71–86. Computational Linguistics, Tilburg University.

Duchier, D. and C. Gardent (2001). Tree Descriptions, Constraints, and Incrementality. In H. Bunt, R. Muskens, and E. Thijsse (Eds.), *Computing Meaning*, pp. 229–246. Kluwer Academic Publishers.

Egg, M., A. Koller, and J. Niehren (2001). The Constraint Language for Lambda Structures. *Journal of Logic, Language and Information 10*, 457–485.

Egg, M., J. Niehren, P. Ruhrberg, and F. Xu (1998). Constraints over lambda structures in semantic underspecification. In *Proc. COLING/ACL*, Montreal CA.

Egg, M. and G. Redeker (2006). Underspecified discourse representation. In A. Benz and P. Kuehlein (Eds.), *Proceedings of the workshop on constraints in discourse*, Ireland. Maynooth University.

van Eijck, J. and H. Kamp (1997). Representing Discourse in Context. In *Handbook of Logic and Language*. Amsterdam: Elsevier.

Erk, K., A. Koller, and J. Niehren (2002). Processing underspecified semantic representations in the constraint language for lambda structures. *Journal of Research on Language and Computation 1*, 127–169.

Forbes, K., E. Miltsakaki, R. Prasad, A. Sarkar, A. K. Joshi, and B. L. Webber (2001). D-ltag system—discourse parsing with a lexicalised tree-adjoining grammar.

Frazier, L. (1995). Issues of Representation in Psycholinguistics. In J. Miller and P.D.Eimas (Eds.), *Speech, Language and Communication*. New York: Academic Press.

Gärdenfors, P. (1988). *Knowledge in Flux, Modeling the Dynamics of Epistemic States*. Cambridge, Massachusetts: MIT Press.

Gardent, C. (1997a). Discourse Tree Adjoining Grammars. CLAUS Report 89, Dept. of Computational Linguistics, University of the Saarland.

Gardent, C. (1997b, January). Parallelism, hou and deaccenting. CLAUS Report 85, Dept. of Computational Linguistics, University of the Saarland.

Gardent, C. and M. Kohlhase (1995, October 24). Focus and Higher Order Colour Unification. Draft.

Gardent, C. and M. Kohlhase (1996). Higher–order coloured unification and natural language semantics. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*. Santa Cruz, USA.

Gardent, C., M. Kohlhase, and N. van Leusen (1996, October). Corrections and Higher Order Unification. In D. Gibbon (Ed.), *Natural Language Processing and Speech Technology.*, Berlin. Mouton de Gruyter.

Gardent, C. and B. Webber (1998). Describing Discourse Semantics. In *Proceedings of the TAG+ workshop*, Philadelphia. University of Pennsylvania.

Geurts, B. (1995). *Presupposing*. Ph. D. thesis, University of Stuttgart.

Geurts, B. (1998). Presuppositions and Anaphors in Attitude Contexts. *Linguistics and Philosophy 21*, 545–601.

Geurts, B. (1999). *Presuppositions and Pronouns*. Oxford: Elsevier.

Geurts, B. (2000). Buoyancy and Strength. *Journal of Semantics 17*, 315–333.

Geurts, B. (2005). Specific indefinites, presupposition, and scope. In R. Bauerle, U. Reyle, and T. E. Zimmermann (Eds.), *Presuppositions and Discourse: essays offered to Hans Kamp.* Amsterdam: Elsevier. Paper available since 2000.

Geurts, B. and R. van der Sandt (1997). Presuppositions and Backgrounds. In P. Dekker, M. Stokhof, and Y. Venema (Eds.), *Proceedings of the 11th Amsterdam Colloquium*, pp. 37–42. ILLC/ University of Amsterdam.

Geurts, B. and R. A. van der Sandt (1998). Domain Restriction. In P. Bosch and R. A. van der Sandt (Eds.), *Focus: Linguistic, Cognitive, and Computational Perspectives*. Cambridge University Press. to appear.

Gilbers, D. and H. de Hoop (1998). Conflicting constraints: An introduction to optimality theory. *Lingua 104*, 1–12.

Ginzburg, J. (1994). An Update Semantics for Dialogue. In H. Bunt, R. Muskens, and G. Rentier (Eds.), *International Workshop on Computational Semantics*, pp. 111–20. ITK, Tilburg University.

Grice, H. (1975). Logic and Conversation. In P. Cole and J. Morgan (Eds.), *Syntax and Semantics 3: Speech Acts*, pp. 41–58. New York: Academic Press. From Grice's William James Lectures, Harvard University, 1967.

Grice, H. (1978). Further Notes on Logic and Conversation. In *Syntax and Semantics*. London: Academic Press, Inc.

Groenendijk, J. and M. Stokhof (1990). Dynamic Montague Grammar. In L. Kálmán and L. Pólos (Eds.), *Papers from the Second Symposium on Logic and Language*, pp. 3–48. Budapest: Akadémiai Kiadó.

Groenendijk, J. and M. Stokhof (1991). Dynamic Predicate Logic. *Linguistics and Philosophy 14*, 39–100.

Grosz, B. and C. Sidner (1986). Attention, Intention, and the Structure of Discourse. *Computational Linguistics 12*, 175–204.

Grosz, B. J., A. K. Joshi, and S. Weinstein (1983). Providing a unified account of definite noun-phrases in discourse. In *Proc. of the 21st Annual Meeting of the ACL*, pp. 44–50. Association of Computational Linguistics.

Grosz, B. J., A. K. Joshi, and S. Weinstein (1995, June). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics 21*(2), 203–225.

Haegeman, L. (1994). *Introduction to Government & Binding Theory* (second edition ed.). Blackwell Publishers.

Hamblin, C. (1971). Mathematical Models of Dialogue. *Theoria 37*, 130–55.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph. D. thesis, Umass, Amherst. published in 1989 by Garland, New York.

Heim, I. (1983). On the Projection Problem for Presuppositions. WCCFL 2.

Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics 9*, 183–221.

Hendriks, P. and H. de Hoop (2001). Optimality Theoretic Semantics. *Linguistics and Philosophy 24*(1), 1–32.

Hirschberg, J. and D. Litman (1993). Empirical studies on the disambiguation of cue phrases. *Computational Linguistics 19*(3), 501–530.

Hobbs, J. (1979). Coherence and Coreference. *Cognitive Science 3*, 67–90.

Hobbs, J. (1985). On the Coherence and Structure of Discourse. CSLI Report.

Huddleston, R. and G. K. Pullum (Eds.) (2002). *The Cambridge Grammar of the English Language*. Cambridge University Press.

Hulstijn, J. (2000). *Dialogue Models for Inquiry and Transaction*. Ph. D. thesis, Universiteit Twente.

Hunter, J. and N. Asher (2005). A presuppositional Account of Indexicals. In P. Dekker and M. Franke (Eds.), *Fiftheenth Amsterdam Colloquium*, pp. 119–124. ILLC/ University of Amsterdam.

Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, Massachusetts: MIT Press.

Johnson, M. (1991). Logic and Feature Structures. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, Sydney, Australia.

Joshi, A. (1987). An introduction to Tree Adjoining Grammars. In A. Manaster-Ramer (Ed.), *Mathematics of Language*, pp. 87–114. Amsterdam: John Benjamins.

Joshi, A. K., L. Levy, and M. Takashaki (1975). Tree Adjunct Grammars. *Journal of Computer and System Science 10*, 136–163.

Joshi, A. K. and Y. Schabes (1991). Tree-Adjoining Grammars and Lexicalised Grammars. In Nivat and Podelski (Eds.), *Definability and Recognizability of Sets of Trees*. Elsevier. See also course material 'Linguistic Aspects of Tree Adjoining Grammars', Fourth ESSLLI Summerschool, Essex, 1992.

Kallmeyer, L. (1999). *Tree Description Grammars and Underspecified Representations*. Ph. D. thesis, University of Tübingen.

Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (Eds.), *Formal Methods in the Study of Language*, pp. 277–322. Amsterdam: Mathematical Centre.

Kamp, H. (2001a). The importance of presupposition. In C. Rohrer and A. Rossdeutscher (Eds.), *Linguistic Form and its Computation*. CSLI Press.

Kamp, H. (2001b). Presupposition computation and presupposition justification: One aspect of the interpretation of multi-sentence discourse. In M. Bras and L. Vieu (Eds.), *Semantics and Pragmatics of Discourse and Dialogue: Experimenting with Current Theories*. Elsevier.

Kamp, H. (2004). Information structure in a dynamic theory of meaning. In *Proceedings of the LSK conference*, Seoul.

Kamp, H. and U. Reyle (1993). *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Number 42 in Studies in Linguistics and Philosophy. Dordrecht: Kluwer Academic Publishers.

Kaplan, J. (1984). Obligatory *too* in English. *Language 60*(3), 510–518.

Karttunen, L. (1969). Pronouns and Variables. In R. B. et al. (Ed.), *Papers from the Fifth Regional Meeting of the Chicago Linguistic Society*, University of Chicago, USA.

Karttunen, L. (1974). Presupposition and Linguistic Context. *Theoretical Linguistics 1*, 181–194.

Keller, F. (1997). Underspecified presuppositions. In *Empirical Issues in Formal Syntax and Semantics: Selected Papers from the Colloque de Syntaxe et de Semantique de Paris (CSSP 1995)*, pp. 291–315. Bern: Peter Lang.

Knott, A. (1996). *A Data-driven Methodology for Motivating a Set of Coherence Relations*. Ph. D. thesis, University of Edinburgh.

Knott, A. and C. Mellish (1996). A feature-based account of the relations signalled by sentence and clause connectives. *Language and Speech 39*(3-2), 143–183.

Knott, A., T. Sanders, and J. Oberlander (Eds.) (2001). *Cognitive Science 12: special issue on Coherence Relations and Connectives*. Elsevier.

Krifka, M. (1993). Focus and Presupposition in Dynamic Interpretation. *Journal of Semantics 10*, 269–300.

Kripke, S. (1979). Speaker's reference and semantic reference. In P. F. et al. (Ed.), *Contemporary perspectives in the Philosophy of Language*. University of Minnesota Press.

Kroch, A. (1987). Unbounded dependencies and subjacency in tree adjoining grammar. In A. Manaster-Ramer (Ed.), *Mathematics of Language*, pp. 143–172. Amsterdam: John Benjamins.

Kroch, A. and A. Joshi (1987). Analysing extraposition in tree adjoining grammar. In G. Huck and A. Ojeda (Eds.), *Syntax and Semantics: Discontinuous Constituency*, pp. 107–149. Academic Press.

van Kuppevelt, J. (1997). Topic and Comment. In Lamarque (Ed.), *Concise Encyclopedia of Philosophy of Language*, pp. 191–198. Elsevier.

Lakoff, R. (1971). If's, and's, and but's about conjunction. In C. Fillmore and D. Langendoen (Eds.), *Studies in Linguistic Semantics*, pp. 114–149. New York: Holt, Rhinehart and Winston.

Lascarides, A. and N. Asher (1991). Discourse relations and defeasible knowledge. In *Proceedings of the 29th Conference of the ACL*, Berkeley, CA, pp. 55–63.

Lascarides, A., N. Asher, and J. Oberlander (1992). Inferring Discourse Relations in Context. In *Proceedings of the 30th Conference of the ACL*, Delaware, pp. 1–8.

van Leusen, N. (2004). Incompatibility: A Diagnosis of Correction. *Journal of Semantics 21*(4), 415–442.

van Leusen, N. and R. Muskens (2003). Construction by Description in Discourse Representation. In J. Peregrin (Ed.), *Meaning, the Dynamic Turn*. Elsevier.

Levinson, S. (1983). *Pragmatics*. Cambridge University Press.

Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.

Lewis, D. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic 8*, 339–359.

Maier, E. (2006). *Belief in Context: Towards a unified semantics of de re and de se attitude reports*. Ph. D. thesis, Radboud Universiteit Nijmegen.

Mann, W. and S. Thompson (1987). Rhetorical Structure Theory: A Theory of Text Organization. ISI Reprint Series:ISI-RS-87-190.

Mann, W. and S. Thompson (1988). Rhetorical structure theory:toward a functional theory of text organization. *Text 8*, 243–281.

Marcu, D. (1997). *The Rhetorical Parsing, Summarization, and Generation of Natural Language Texts*. Ph. D. thesis, Dept. of Computer Science, University of Toronto.

Marcu, D. (2000). The Rhetorical Parsing of Unrestricted Texts: a Surface-based Approach. *Computational Linguistics 26*(3), 395–448.

Marcus, M., D. Hindle, and M. Fleck (1983). D-theory: Talking about Talking about Trees. In *Proceedings of the 21st ACL*, pp. 129–136.

Muskens, R. (1995). Order-independence and Underspecification. DYANA Report R2.2.C, University of Amsterdam.

Muskens, R. (1996). Combining Montague Semantics and Discourse Representation. *Linguistics and Philosophy 19*, 143–186.

Muskens, R. (2001). Talking about Trees and Truth-Conditions. *Journal of Logic, Language and Information 10*, 417–455.

Muskens, R. and E. Krahmer (1998). Description Theory, LTAGs and Underspecified Semantics. In *Fourth International Workshop on Tree Adjoining Grammars and Related Frameworks*, Philadelphia, PA, pp. 112–115. Institute for Research in Cognitive Science.

Muskens, R., J. van Benthem, and A. Visser (1997). Dynamics. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 587–648. Amsterdam: Elsevier.

Niehren, J., M. Pinkal, and P. Ruhrberg (1997). A uniform approach to underspecification and parallelism. In *Proceedings ACL'97*, Madrid, pp. 410–417.

Partee, B. H. (1991). Topic, Focus and Quantification. In *Cornell Working Papers in Linguistics*, Volume 10. SALT 1.

Partee, B. H., A. ter Meulen, and R. E. Wall (1990). *Mathematical Methods in Linguistics*. Studies in Linguistics and Philosophy. Kluwer Academic Publishers.

Pierrehumbert, J. and J. Hirschberg (1990). The Meaning of Intonation Contours in the Interpretation of Discourse. In Cohen, Morgan, and Pollack (Eds.), *Intentions in Communication*, pp. 271–323. Cambridge: MIT Press.

Pinkal, M. (1996). Radical Underspecification. In *Proceedings of the Tenth Amsterdam Colloquium*, pp. 587–606.

Polanyi, L. (1985). A Theory of Discourse Structure and Discourse Coherence. In *Papers from the general session of the Chicago Linguistic Society, CLS 21*, pp. 306–322.

Polanyi, L. (1988). A formal model of discourse structure. *Journal of Pragmatics 12*, 601–638.

Polanyi, L., C. Culy, M. van den Berg, G. L. Thione, and D. Ahn (2004). Sentential structure and discourse parsing. In *Proceedings of the ACL workshop on ... 2004*.

Polanyi, L. and M. H. van den Berg (1996). Discourse structure and discourse interpretation. In P.Dekker and M.Stokhof (Eds.), *Proceedings of the Tenth Amsterdam Colloquium*, Volume 1, pp. 113–131. University of Amsterdam.

Prince, A. and P. Smolensky (1997). Optimality theory: From neural networks to universal grammar. *Science 275*, 1604–1610.

Prüst, H. (1992). *On Discourse Structuring, VP Anaphora, and Gapping*. Ph. D. thesis, University of Amsterdam.

Prüst, H., R. Scha, and M. van den Berg (1994). Discourse grammar and verb phrase ananphora. *Linguistics and Philosophy 17*, 261–327.

Pulman, S. (1997). Higher Order Unification and the Interpretation of Focus. *Linguistics and Philosophy 20*, 73–115.

Rambow, O. (1994a). *Formal and Computational Aspects of Natural language Syntax*. Ph. D. thesis, University of Pennsylvania.

Rambow, O. (1994b). Multiset-valued linear index grammars: Imposing dominance constraints on derivations. In *Proceedings of ACL*.

Rambow, O., K. Vijay-Shanker, and D. Weir (1995a). D-Tree Grammars. In *Proceedings of ACL*.

Rambow, O., K. Vijay-Shanker, and D. Weir (1995b). Formal properties of d-tree grammars. In *Proceedings of the Fourth Workshop on the Mathematics of Language*.

Reinhart, T. (1976). *The syntactic domain of anaphora*. Ph. D. thesis, MIT, Cambridge.

Rescher, N. (1976). *Plausible Reasoning: an introduction to the theory and practice of plausibilistic inference.* Assen/ Amsterdam: Van Gorcum.

Reyle, U. (1993). Dealing with ambiguities by underspecification: Construction, representation and deduction. *Journal of Semantics 10*, 123–179.

Riester, A. (2005). Context and focus projection. a compositional, intonation-based account of focus interpretation. In E. Maier, C. Bary, and J. Huitink (Eds.), *Proceedings of Sinn und Bedeutung 9*.

Rogers, J. (1994). *Studies in the Logic of Trees with Applications to Grammar Formalisms*. Ph. D. thesis, University of Delaware.

Rogers, J. and K. Vijay-Shanker (1992). Reasoning with Descriptions of Trees. In *Proceedings of ACL*.

Rogers, J. and K. Vijay-Shanker (1994). Obtaining trees from their descriptions: an application to tree adjoining grammars. *Computational Intelligence 10*(4), 401–421.

van Rooij, R. (2005). Presupposition: An (un)common attitude? In R. Bauerle, U. Reyle, and T. E. Zimmermann (Eds.), *Presuppositions and Discourse: essays offered to Hans Kamp.* Amsterdam: Elsevier.

Rooth, M. (1985). *Association with Focus.* Ph. D. thesis, University of Massachusetts, Amherst.

Rooth, M. (1992). A Theory of Focus Interpretation. *Natural Language Semantics 1*, 75–116.

van der Sandt, R. (1988). *Context and Presupposition.* London: Croon Helm.

van der Sandt, R. (1989). Presupposition and Discourse Structure. In R. Bartsch, J. van Benthem, and P. van Emde Boas (Eds.), *Semantics and Contextual Expression*, pp. 259–296. Dordrecht: Foris.

van der Sandt, R. (1991). Denial. In *Papers from the annual regional meeting of the Chicago Linguistic Society*, Volume 27, pp. 331–344.

van der Sandt, R. (1992). Presupposition Projection as Anaphora Resolution. *Journal of Semantics 9*, 333–377.

van der Sandt, R. (2007). Processing Descriptions. *Manuscrito – Rev.Int.Fil. 29.* To appear.

van der Sandt, R. and B. Geurts (1991). Presupposition, anaphora, and lexical content. In O. Herzog and C.-R. Rollinger (Eds.), *Text Understanding in LiLOG*, pp. 259–296. Springer, Berlin.

van der Sandt, R. and E. Maier (2003). Denials in Discourse. Abstract.

Scha, R. and L. Polanyi (1988). An augmented context-free grammar for discourse. In *Proceedings of the 12th international Conference on Computational Linguistics (COLING)*, Budapest, pp. 22–27.

Schabes, Y. (1990). *Mathematical and Computational Aspects of Lexicalised Grammars.* Ph. D. thesis, University of Pennsylvania.

Schilder, F. (1997). Tree discourse grammar, or how to get attached to a discourse. In *Proceedings of the Tilburg Conference on Formal Semantics*, Tilburg, Netherlands.

Searle, J. (1969). *Speech Acts.* Cambridge: Cambridge University Press.

Shieber, S. (1985). Evidence against the context-freeness of natural language. *Linguistics and Philosophy 8*, 333–343. Reprinted in Savitch et al. (1987).

Sidner, C. (1986). Focusing in the Comprehension of Definite Anaphora. In B. J. Grosz, K. Jones, and B. L. Webber (Eds.), *Readings in Natural Language Processing*, pp. 363–394. Morgan Kaufmann.

Sporleder, C. and A. Lascarides (2007). Using Automatically Labelled Examples to Classify Rhetorical Relations: An Assessment. *Natural Language Engineering*. To appear.

Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics*, Volume 9, pp. 315–332. New York: Academic Press.

Stalnaker, R. (1998). On the Representation of Context. *Journal of Logic, Language and Information 7*, 3–19.

Stalnaker, R. C. (1974). Pragmatic Presuppositions. In M. Munitz and P. Unger (Eds.), *Semantics and Philosophy*, pp. 197 – 213. New York: New York University Press.

Stone, M. and B. Webber (1998). Textual Economy through Closely Coupled Syntax and Semantics. In *Proceedings of the Ninth International Workshop on Natural Language Generation*, Niagara-on-the-Lake, Canada, pp. 178–187.

Thompson, W. (2002). Combining Higher-Order Unification with Syntactic Underspecification. In *Lola 7 ??*, Pecs, Hungary, pp. 151–161.

Vallduví, E. (1990). *The Informational Component*. Ph. D. thesis, University of Pennsylvania.

Veltman, F. (1996). Defaults in Update Semantics. *Journal of Philosophical Logic 25*(3), 221–261. An older version appeared as ILLC report LC-91-02, 1991, University of Amsterdam.

Vijay-Shanker, K. (1992). Using descriptions of trees in a tree adjoining grammar. *Computational Linguistics 18*, 481–518.

Vijay-Shanker, K. and D. Weir (1994). The equivalence of four extensions of context-free grammar. *Mathematical Systems Theory 27*, 511–545.

Visser, A. (1994). The Design of Dynamic Discourse Denotations. Lecture Notes, Utrecht University.

Von Stechow, A. (1991). Current issues in the theory of focus. In V. Stechow and Wunderlich (Eds.), *Semantik: ein Internationales Handbuch der Zeitgenössische Forschung*. Berlin, de Gruyter.

Walker, M. (1996a). Inferring Acceptance and Rejection in Dialogue by Default Rules of Inference. *Language and Speech 39*, 265–304.

Walker, M. A. (1996b). Limited Attention and Discourse Structure. *Computational Linguistics 22*(2), 255–264.

Webber, B. (1978). *A Formal Approach to Discourse Anaphora*. Ph. D. thesis, Harvard University.

Webber, B. (1991). Structure and ostension in the interpretation of discourse deixis. *Language and Cognitive Processes 6*(2), 107–135.

Webber, B. (2004). D-LTAG: extending lexicalised TAG to discourse. *Cognitive Science 28*, 751–779.

Webber, B. and A. Joshi (1998). Anchoring a lexicalised tree-adjoining grammar for discourse. In *Coling/ACL Workshop on Discourse Relations and Discourse Markers*, Montreal, Canada, pp. 86–92.

Webber, B., A. Joshi, M. Stone, and A. Knott (2003). Anaphora and Discourse Structure. *Computational Linguistics 29*(4), 545–587.

Webber, B., A. Knott, and A. Joshi (1999). Multiple Discourse Connectives in a Lexicalized Grammar for Discourse. In H. Bunt and E. Thijsse (Eds.), *Proceedings of the 3rd International Workshop on Computational Semantics (IWCS-3)*, Tilburg, pp. 309–325.

Webber, B., A. Knott, M. Stone, and A. Joshi (1999). Discourse relations: A structural and presuppositional account using lexicalised tag. In *Proceedings of the ACL'99*.

Weir, D. (1988). *Characterising mildly context-sensitive grammar formalisms*. Ph. D. thesis, University of Pennsylvania.

Zeevat, H. (1989). A compositional approach to discourse representation theory. *Linguistics and Philosophy 12*, 95–131.

Zeevat, H. (1992). Presupposition and Accommodation in Update Semantics. *Journal of Semantics 9*, 379–412.

Zeevat, H. (1999). Old material. In B. Geurts, M. Krifka, and R. van der Sandt (Eds.), *Focus and Presupposition in Multi-Speaker Discourse*, Workshop reader. ESSLI.

Zwarts, F. (1995). Nonveridical contexts. *Linguistic Analysis 25*(3–4), 286–312.

Summary

This work presents the *Logical Description Grammar framework of discourse interpretation*, and applies the model in the domain of anaphora resolution and presupposition theory. The purpose of the enterprise is to provide what might be called a common language for the specification of discourse theory. Much of current research in linguistics and discourse semantics focuses on highly specialised issues. Moreover, it employs a number of different formal representation systems, so that the analyses provided are sometimes hard to compare. Although it is unavoidable that alternative treatments of one and the same topic arise in a field of research where we may meet with uncertainties of various kinds, such a dispersion of linguistic findings need not be accepted. Even when we concentrate on details or isolated phenomena, linguistic theory is essentially a theory of the *whole* of a language user's linguistic capabilities. Analyses of different topics must be combined so that they may help to evaluate and strengthen each other.

The LDG framework of discourse interpretation is a flexible system which allows much freedom in setting the parameters of discourse theory. As such it may help to integrate and compare findings currently spread over different approaches or spelled out in different formal theories. Thus, a shared discourse theoretical perspective may be developed.

The LDG framework is a *reasoning system*, the central component of which is a so-called Logical Description Grammar (LDG) for discourse. The grammar represents the language user's linguistic knowledge. To obtain the full meaning of a discourse, however, linguistic knowledge is not enough. Language users employ world knowledge and knowledge of what is normally the case as well. They are able to infer a most plausible interpretation of a discourse, given their combined linguistic and nonlinguistic knowledge. Therefore, the framework of discourse interpretation also includes a reservoir of world knowledge, and a so-called default reasoning component and preference system.

A standard (non-LDG) discourse grammar generates linguistic representations or discourse representations for a given discourse. Usually, these are *tree struc-*

tures, decorated with syntactic labels and features, lexemes, and semantic values. If a single sentence or discourse has more than one possible analysis or interpretation, the grammar generates a different representation for each of them. We say, in that case, that the discourse is *ambiguous* or *underspecified*. An example is given in (B.1); this is a common case of underspecification at discourse level, where the relation of coherence linking two statements is underspecified.

(B.1) *John fell. Bill pushed him.*

We may fill in the implicit relation between the events described in (B.1) differently. Either John fell, and subsequently, Bill pushed him. Or we infer that John fell *because* Bill pushed him, so the pushing took place before the falling.

Ambiguity and underspecification may come in various forms and guises in a given language, and discourse abounds with it. A grammar formalism can capture the interpretation of an ambiguous sentence or discourse by producing all possible linguistic representations and corresponding meanings. However, it can also produce a single representation, which underspecifies the syntactic or semantic properties that are variable and differ across the possible representations. *Description grammars* are of the latter kind. For any given sentence or discourse they generate a single description, which specifies all that is invariable about the input. In (B.1), it will say that there is a relation between the events expressed by the two sentences, but it will not specify which one.

The LDG for discourse employed here is developed in Van Leusen and Muskens (2003) building on Vijay-Shanker (1992), Muskens (1995; 2001), and Gardent and Webber (1998). It presupposes two levels of linguistic analysis, namely that of the *descriptions*, and that of the *linguistic representations* they describe. Discourse is processed incrementally. Each processing step results in an updated *discourse description*. This is a set of statements in a logical language (type logic) representing the language user's knowledge of the discourse. The grammar is a set of statements in type logic as well. The combined theory of the grammar and a discourse description describes the syntax and semantics of the discourse in parallel. Since descriptions can be partial, syntactic or semantic properties of a discourse may remain underspecified. For any given discourse description, its linguistic representations can be obtained by reasoning about the tree structures and semantic values that 'fit' or 'verify' the description, given his grammar. Discourse *parsing* comes down to logical deduction.

Sentence grammar is fully integrated with discourse grammar. As in Lexicalised Tree Adjoining Grammar (Joshi 1987; Joshi and Schabes 1991) the lexicon is the 'motor' of the grammar system. All the nodes, structure, and semantic values occurring in a discourse tree are contributed by lexemes or implicit discourse relations described in the lexicon, and occurring in the discourse. Beside a unified treatment of sentence and discourse analysis, the LDG architecture guarantees a smooth integration of grammatical levels. The grammar describes syntax and

semantics in tamdem, and the preference system covers pragmatics aspects of interpretation. The language user may access all knowledge sources in parallel.

In order to provide a full-fledged, context-sensitive compositional semantics, discourse tree structures are decorated with various kinds of semantic information. Each node in a discourse tree carries a *semantic value*, a *local context* composed in the style of Karttunen (1974), and a *background value*. The semantic representation language is fine-grained CDRT, a variant of compositional DRT put forward in Van Leusen and Muskens (2003). The semantic value of the root is constructed compositionally from the semantic values and background values of the nodes it dominates. Furthermore, each discourse comes with a underspecified general background value B, which represents what is implicitly taken for granted by the participants in the conversation. *Discourse meaning* is defined as the update of this implicit background B with the semantic value of the root node of the discourse tree.

The framework of discourse interpretation is applied in the domain of anaphoricity and presupposition theory. It is shown that with a discourse semantics that is sufficiently fine-grained and rich, approaches to presupposition projection as different in perspective as the so-called *satisfaction approach* to presupposition (Karttunen 1974; Heim 1983; Beaver 1995) and the *anaphoric approach* (Van der Sandt 1992; Geurts 1999) can be integrated in a single treatment. The first treats presuppositions as tests on the context of interpretation. For a sentence to be felicitous or interpretable, the context of interpretation must entail its presuppositions. Sometimes a presupposition can only be satisfied by strengthening the context of interpretation; this is 'accommodation' in the sense of the satisfaction approach. The central claim of the second approach is that presuppositional material must be bound in the discourse context of the trigger. If it cannot be bound, it must be inserted at some semantically accessible location in the discourse context; this is 'accommodation' in the sense of the anaphoric approach.

In the LDG framework, anaphoric elements introduce underspecified discourse markers. Through properness conditions on local contexts and the discourse meaning, the grammar enforces that these discourse markers are bound. Anaphora resolution is a reasoning process, in which a language user determines possible binders for a given anaphoric marker and selects solutions that result in a most preferred reading of the discourse. Presuppositional elements are anaphoric, but they also introduce a condition that requires the entailment of the descriptive material they carry in their local context. By letting presupposition triggers introduce presuppositional conditions while at the same time analysing them as anaphors, the gist of both of the above mentioned approaches is captured. The relations of *properness* and *entailment* are made available at object level in fine-grained CDRT.

The resulting fused presupposition theory is part of the discourse theory: it is formulated exclusively in terms of the interaction of general parameters of the discourse theory with the syntactic-semantic characterisation of presuppo-

sition triggers in the lexicon. Discourse parameters that play a central role are *anaphoricity*, and the *givenness* or inferrability of information in the context of interpretation. A third factor is *accommodation*, that is, the specification of background information when the interpretation task calls for it. A fourth one is the constraining effect of *preferences*. Main preferences in this treatment concern the consistency and informativity of discourse units relative to their local context, the plausibility of resulting readings, accommodation at top level in the discourse meaning, and 'economy'—which applies to the measure to which readings require accommodation.

Interestingly, accommodation comes in two forms, which can co-occur. One is 'global accommodation', and consists in the specification of the implicit general background of the discourse; in line with the satisfaction approach common knowledge may freely be accepted in B—though subject to consistency constraints and the like. The other is 'content modification', and consists in the insertion of background material in locations in the discourse meaning that are accessible from the presupposition trigger. In line with the anaphoric approach what is accommodated is precisely the presuppositional material, including binders.

The predictions of the resulting presupposition theory obviously depend on the setting of preferences, and choices in the implementation of the discourse theory in general. When presuppositional elements are analysed as anaphors, the predictions conform to those of the anaphoric approach, though generally speaking more readings will be distinguished. This is due to the fact that the treatment of accommodation is richer, and more of the account is handled with preferences. What comes out as a most preferred reading is the result of the interaction of all grammatical constraints and preferences involved. Thus, sentences that may be said to evoke weak presuppositions can get strongly presupposing interpretations. In terms of explanatory force the satisfaction approach and the anaphoric approach to presupposition can be seen to strengthen each other. This is particularly clear with accommodation, where 'satisfaction type' global accommodation may co-occur with 'anaphoric style' content modification. Apart from this, the presuppositional condition functions as a suitability constraint on potential binders of the presuppositional element.

Some spin-off of the implementation in LDG, as in other underspecification formalisms, is that a transparant declarative treatment is provided of aspects of presupposition projection that are often handled procedurally in the anaphoric approach to presupposition. Furthermore, because of the flexibility and generality of the analysis it can be expected to cover a wide range of presupposition triggers, including bridging anaphors. Finally, some insight in the role of local informativity and consistency conditions in discourse theory is gained. It is tentatively concluded that these can better be handled by means of the preference system or soft constraints, rather than hard constraints.

Curriculum Vitae

Noor van Leusen was born in 1967 in Arnhem. She studied Language and Literature at Tilburg University, starting in 1985 and graduating *cum laude* in 1991. She worked at the department of Alfa-Informatica of the University of Amsterdam from 1991 to 1995, and continued her research with the ILLC in Amsterdam till 2003. Her research topics are the semantics of corrections and denials in dialogue, the interaction with focus interpretation, and discourse theory in general.

From 1996 to 1997 she worked at the department of computational linguistics of Saarbrücken University, investigating corrective update from the perspective of Higher Order Unification theory. In 1998/1999 she was a visitor at the IMS in Stuttgart, joining the research group on lexical inference and underspecified meaning representation. From 1999 to 2005 linguistic research was interspersed with part-time work at Nijmegen University Library and Ticer BV.

In 2004 she joined the philosophy of language department of Nijmegen University. Supervised by Prof. dr. Rob van der Sandt she rounded off her Ph.D., focusing her research on description theory for discourse and presupposition theory. She published work on discourse theory, the semantics of corrections, focus interpretation, and presupposition theory.