

The Semantics and Pragmatics of Presupposition

Nicholas Asher

Department of Philosophy,
University of Texas at Austin,
Austin, Texas 78712,
USA
nasher@bertie.la.utexas.edu

Alex Lascarides

Centre for Cognitive Science and HCRC,
University of Edinburgh,
2, Buccleuch Place,
Edinburgh, EH8 9LW,
Scotland, UK
alex@cogsci.ed.ac.uk

August 12, 1998

Abstract

In this paper, we offer a novel analysis of presuppositions, paying particular attention to the interaction between the knowledge resources that are required to interpret them. The analysis has two main features. First, we capture an analogy between presuppositions, anaphora and scope ambiguity (cf. van der Sandt, 1992), by utilising semantic underspecification (cf. Reyle, 1993). Second, resolving this underspecification requires reasoning about how the presupposition is rhetorically connected to the discourse context.

This has several consequences. First, since pragmatic information plays a role in computing the rhetorical relation, it also constrains the interpretation of presuppositions. Our account therefore provides a formal framework for analysing problematic data, which require pragmatic reasoning. Second, binding presuppositions to the context via rhetorical links replaces accommodating them, in the sense of adding them to the context (cf. Lewis, 1979). The treatment of presupposition is thus generalized and integrated into the discourse update procedure.

We formalise this approach in *SDRT* (Asher 1993, Lascarides and Asher 1993), and demonstrate that it provides a rich framework for interpreting presuppositions, where semantic and pragmatic constraints are integrated.

1 Introduction

The interpretation of a presupposition typically depends on the context in which it is made. Consider for instance, sentences (1) vs (2), adapted from van der Sandt (1992); the presupposition triggered by *Jack's son* (that Jack has a son) is implied by (1), but not by (2).

- (1) If baldness is hereditary, then Jack's son is bald.
- (2) If Jack has a son, then Jack's son is bald.

The challenge for a formal semantic theory of presuppositions is to capture contextual effects such as these in an adequate manner. In particular, such a theory must account for why

the presupposition in (1) projects from an embedded context, while the presupposition in (2) does not. This is a special case of the Projection Problem: If a compound sentence S is made up of constituent sentences S_1, \dots, S_n , each with presuppositions P_1, \dots, P_n , then what are the presuppositions of S ?

Many recent accounts of presupposition that offer solutions to the Projection Problem have exploited the dynamics in dynamic semantics (e.g., Beaver 1996, Geurts 1996, Heim 1982, van der Sandt 1992). In these frameworks, assertional meaning is a relation between an input context (or information state) and an output context. Presuppositions impose tests on the input context, which researchers have analyzed in two ways: either the context must satisfy the presuppositions of the clause being interpreted (e.g., Beaver 1996, Heim 1982) or the presuppositions are anaphoric (e.g., van der Sandt, 1992) and so must be bound to elements in the context. But clauses carrying presuppositions can be felicitous even when the context fails these tests (e.g., (1)). A special purpose procedure known as *accommodation* is used to account for this (cf. Lewis, 1979): if the context fails the presupposition test, then the presupposition is accommodated or *added* to it, provided various constraints are met (e.g., the result must be satisfiable).

This combination of test and accommodation determines the projection of a presupposition. For example in (1), the antecedent produces a context which fails the test imposed by the presupposition in the consequent (either satisfaction or binding). So it's accommodated. Since it can be added to the context outside the scope of the conditional, it can project out from its embedding. In contrast, the antecedent in (2) ensures that the input context passes the presupposition test. So the presupposition isn't accommodated, the input context isn't changed, and the presupposition isn't projected out from the conditional.

Despite these successes, this approach has trouble with some simple predictions. Compare the following two dialogues (3abc) and (3abd):

- (3)
- a. A: Did you hear about John?
 - b. B: No, what?
 - c. A: He had an accident. A car hit him.
 - d. A: He had an accident. ??The car hit him.

The classic approach we just outlined would predict no difference between these two discourses and would find them both acceptable. But (3abd) is unacceptable. As it stands it lacks discourse coherence, while (3abc) does not; the presupposition of *the car* cannot be accommodated in (3abd). We will argue that the proper treatment of presuppositions in discourse, like a proper treatment of assertions, requires a notion of discourse coherence and must take into account the rhetorical function of both presupposed and asserted information. We will provide a formal account of presuppositions, which integrates constraints from compositional semantics and pragmatics in the required manner.

We will start by examining van der Sandt's theory of presupposition satisfaction, since he offers the most detailed proposal concerning accommodation. We will highlight some difficulties, and offer a new proposal which attempts to overcome them. We will adopt van der Sandt's view that presuppositions are anaphoric, but give it some new twists. First, like

other anaphoric expressions (e.g., anaphoric pronouns), presuppositions have an *underspecified* semantic content. Interpreting them in context involves resolving the underspecification. The second distinctive feature is the way we resolve underspecification. We assume a formal model of discourse semantics known as SDRT (e.g., Asher 1993, Lascarides and Asher 1993), where semantic underspecification in a proposition is resolved by reasoning about the way that proposition rhetorically connects to the discourse context. Thus, interpreting presuppositions becomes a part of discourse update in SDRT.

This has three important consequences. The first concerns pragmatics. SDRT provides an explicit formal account of how semantic and pragmatic information interact when computing a rhetorical link between a proposition and its discourse context. This interaction will define the interpretation of presuppositions, and thus provide a richer source of constraints on presuppositions than standard accounts. This account of presuppositions will exploit pragmatic information over and above the clausal implicatures of the kind used in Gazdar's (1979) theory of presuppositions. We'll argue in §2 that going beyond these implicatures is necessary to account for some of the data.

The second consequence of interpreting presuppositions in SDRT concerns accommodation. In all previous dynamic theories of presupposition, accommodation amounts to adding, but not relating, the presupposed content to some accessible part of the context. This mechanism is peculiar to presuppositions; it doesn't feature in accounts of any other phenomena, including other anaphoric phenomena. In contrast, we model presuppositions entirely in terms of the SDRT discourse update procedure. We replace the notion that presuppositions are added to the discourse context with the notion that they are rhetorically linked to it. Given that the theory of rhetorical structure in SDRT is used to model a wide range of linguistic phenomena when applied to assertions, it would be odd if presupposed information were to be entirely insensitive to rhetorical function. We will show that presupposed information is sensitive to rhetorical function and that the notion of accommodation should be replaced with a more constrained notion of discourse update.

The third consequence concerns the compositional treatment of presupposition. Our approach affords what one could call a compositional treatment of presuppositions. The discourse semantics of SDRT is compositional upon discourse structure: the meaning of a discourse is a function of the meaning of its parts and how they are related to each other. In SDRT presuppositions, like assertions, generate underspecified but interpretable logical forms. The procedure for constructing the semantic representation of discourse takes these underspecified logical forms, resolves some of the underspecifications and relates them together by means of discourse relations representing their rhetorical function in the discourse. So presuppositions have a content that contributes to the content of the discourse as a whole. Indeed, presuppositions have no less a compositional treatment than assertions.

Our discourse based approach affords a wider perspective on presuppositions. Present dynamic accounts of presupposition have concentrated on phenomena like the Projection Problem. For us the Projection Problem amounts to an important special case, which applies to single sentence discourses, of the more general 'discourse' problem: how do presuppositions triggered by elements of a multi-sentence discourse affect its structure and content? We aim to tackle this question here. And we claim that a rich notion of discourse structure, which utilises rhetorical relations, is needed.

While we believe that our discourse based theory of presupposition is novel, we hasten to add that many authors on presupposition like Beaver (1995) and van der Sandt (1992) would agree with us that the treatment of presupposition must be integrated with a richer notion of discourse structure and discourse update than is available in standard dynamic semantics (e.g. Kamp and Reyle's DRT, Dynamic Predicate Logic or Update Semantics), because they believe that pragmatic information constrains the interpretation of presuppositions. We wish to extend their theories with this requisite notion of discourse structure.

2 Van der Sandt's Dynamic Account and its Problems

Van der Sandt (1992) views presuppositions as anaphors with semantic content. He develops this view within the framework of DRT (Kamp and Reyle, 1993), in order to exploit its constraints on anaphoric antecedents.

A presupposition can bind to an antecedent only if there's the same content in either an accessible part of the DRS which represents the discourse context, or an accessible part of the DRS which represents the current clause (i.e., the clause that introduced the presupposition trigger). In (2), for example, the antecedent of the conditional is accessible to the consequent, and it contains the same content as the presupposition that's triggered there.

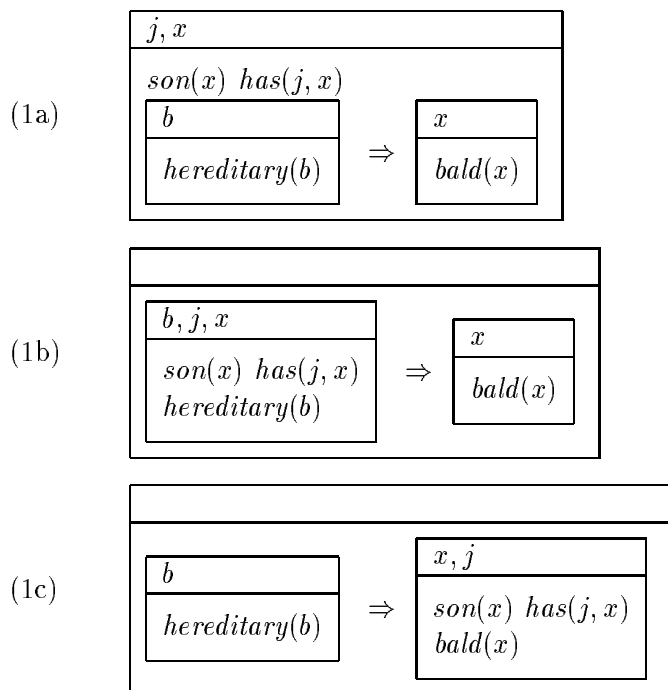
- (1) If baldness is hereditary, then Jack's son is bald.
- (2) If Jack has a son, then Jack's son is bald.

So this presupposition binds to it. This provides a representation of (2) which can be paraphrased as *If Jack has a son, then he is bald*, which matches intuitions.

In contrast, the presupposition in (1) can't be bound, because the context lacks the required content. Following Karttunen (1974) and Heim (1982), van der Sandt (1992) resorts to accommodation: he adds the presupposition to the context (cf. Stalnaker 1974, Lewis, 1979).

Van der Sandt provides an algorithm which specifies how binding and accommodation jointly model presupposition satisfaction. First, the presupposed material is separated from the asserted material in the DRS which represents the current sentence (which may be complex in that it contains several clauses), in order to allow them to be processed differently. One handles the presupposed material first. If it can be bound in the manner specified above, then it is. Otherwise, it's added to an accessible site. One then adds the DRS which represents the asserted material of the current sentence to the DRS representing the previous sentences in the discourse, or some subDRS of it, via DRT's notion of update (note that one of these DRSS may have been modified with the addition of the presupposition). Essentially, DRT's notion of update is set union on both the discourse referents and the DRS conditions.

When the contexts are structurally complex (i.e. contain subDRSS and complex conditions), different possibilities for accommodation arise. Van der Sandt distinguishes between *global accommodation* (as in (1a)), *intermediate accommodation* (as in (1b)) and *local accommodation* (as in (1c)).



These capture the different possible inferences that presuppositions can give rise to. (1a) means: Jack has a son, and if baldness is hereditary, then Jack's son is bald. (1b) means: If baldness is hereditary and Jack has a son, then Jack's son is bald. (1c) means: If baldness is hereditary, then Jack has a son and he is bald. Only (1a) reflects the case where the presupposition projects from the embedding.

Accommodation is subject to certain constraints: the result of the addition should be logically consistent and should not render any part of the asserted content uninformative. Furthermore, van der Sandt argues that if these constraints yield a choice as to where to accommodate the presupposition, then one prefers to add it to the most superordinate or highest DRS context in which the constraints are satisfied. So in (1), global accommodation (i.e., (1a)) is predicted.

Van der Sandt's account of presupposition is compelling, because he offers a precise solution to the Projection Problem. A presupposition projects from an embedded context only if the above algorithm predicts that it is accommodated at a superordinate site in the DRS (e.g., (1a)). There are, however, a number of difficulties with the theory's predictions, some of which are particular to van der Sandt's formulation of accommodation, others of which are endemic to the *semantic* notion of accommodation itself—which we will take henceforth to mean the addition of presupposed information to the context.

The first problem, which is particular to van der Sandt's account, is that local accommodation isn't predicted in certain cases when it should be. Consider (4) (modified from Beaver (1997)):

- (4) Either John didn't solve the problem or else Mary realises that the problem has been solved.

The second disjunct presupposes that the problem has been solved. There's no suitable

accessible antecedent, and so this information has to be accommodated. (4) is reminiscent of one discussed by van der Sandt (1992):¹

- (5) Either John has no donkey or his donkey is eating quietly in the stable.

The presupposition in (5) that John has a donkey cannot be accommodated globally in his theory, because this would render the first disjunct of (5) uninformative: assuming ϕ is the proposition that John has a donkey, then accommodating globally would produce a semantic representation of the form $\phi \wedge (\neg\phi \vee \psi)$, which is equivalent to $\phi \wedge \psi$. So by appealing to the informativeness constraint we mentioned above, van der Sandt predicts that global accommodation is not possible for (5).²

We can't block global accommodation in (4) with the informativeness constraint, however. Adding the presupposition that someone solved the problem doesn't render the first disjunct, that *John* didn't solve the problem, uninformative. Since this global accommodation also results in a consistent discourse, van der Sandt's representation of (4) amounts to: The problem has been solved, and either John didn't solve it or Mary realizes that it has been solved. As Beaver points out, this is contrary to intuitions, which favor local accommodation: Either John didn't solve the problem, or it's been solved and Mary realizes it's been solved.

To capture the intuitive semantics of sentence (1), Van der Sandt relies on the preference for global accommodation. But in (4) this preference causes problems, because local accommodation should be preferred, and yet the constraints on accommodation are satisfied at the superordinate site. We'll see in the course of this section that this puzzle arises because of a general observation about presuppositions that DRT alone can't capture: the preferred site for accommodation depends on a variety of pragmatic factors (cf. Beaver, 1996) that aren't represented in DRSS.

Nevertheless, (4) is in fact a bit odd unless it's uttered in a particular context; we need to know which problem is being talked about. Van der Sandt's theory fails to capture this as well, because it predicts that the presupposition that there is a problem, which is triggered by *the problem*, can be globally accommodated. But discourse contexts which improve the acceptability of (4) strongly suggest that the way presuppositions are satisfied depends on the rhetorical structure of the discourse as a whole. Dialogue (6) illustrates this:

- (6) a. A: The Problem Solving Group is given a problem each day, and the group leader Mary has to assign it to someone in her group. John is the best problem solver. But when he solves a problem, he always boasts about it. This annoys Mary, and so if she thinks that the day's problem is an unsolved one, she gives it to him, to test him. Otherwise, she gives it to someone else.
- b. B: John's being very quiet just now. Did she give him today's problem?
- c. A: Well, I'm not sure she did. Either John didn't solve the problem or else Mary realized that the problem's been solved.

¹We are indebted to a reviewer from *Journal of Semantics* for pointing this out to us.

²We could make this notion of informativeness more precise in a way proposed by Michael Morreau: Any accommodation of a presupposition ϕ in a context C of a sentence ψ must not make ψ in C logically equivalent to χ , such that $\chi \models \phi$ but it's not the case that $\phi \models \chi$.

B's question identifies the problem being talked about, and this improves the acceptability of (4). In addition, the rhetorical links between the (6a), (6b) and (6c) help one interpret the presupposition that the problem has been solved.

But our more basic problem remains: to explain the local accommodation of the presupposition of *realize* in (6c). We could with Geurts (1995) integrate in Van der Sandt's theory further constraints on accommodation such as those in Gazdar's theory, which are based on Grice's notion of conversational implicature. Gazdar's theory predicts local accommodation in any context where by uttering the sentence, the speaker implicates that the presupposition is *not* taken for granted (i.e., the speaker implicates that he believes that it is possible that the presupposition is false). Would such a constraint block global accommodation in (6)? *B*'s question about John's behavior in (6b) serves as an elaboration to the generalization about who gets problems in (6a), and so the answer to the question hangs on whether Mary believes that the problem is solved or not. *A* implies by his utterance (6c) that he doesn't know whether Mary gave John the problem. But what implicatures does this raise? Given what *A* said in (6a), *A* and *B* would know the answer to the question as to whether John got today's problem, if they knew if Mary believed it was solved. So this context implicates that *A* and *B* don't know if Mary believes the problems been solved. However, the context does *not* implicate that *A* and *B* themselves believe the presupposition—that the problem's been solved—is in doubt; the context is mute on this issue. So Gazdar's constraints don't predict local accommodation in this case. Something more is needed to explain why local accommodation occurs here.

To predict local accommodation in (6), we analyze presuppositions in terms of the rhetorical structure of the discourse context. Rhetorical structure is something about which DRT has little to say. We claim that rhetorical structure explains why local accommodation is preferred in some cases where global accommodation would be informative and consistent, and furthermore would meet Gazdar's constraints on how Gricean-style conversational implicatures determine global vs. local accommodation. The place where presuppositions get accommodated depends on the rhetorical links between propositions in discourse as well as their content. In particular, presuppositions are interpreted so that the rhetorical links are as strong as possible. We will return to the detailed analysis of (4) in §5.3 and demonstrate that the *Contrast* relation between the disjuncts in (4) is at its strongest in (6) if the presupposition is accommodated locally.

The difference in acceptability between (3c-d) also demonstrates that the constraints of informativeness and consistency are too weak:

- (3)
- a. A: Did you hear about John?
 - b. B: No, what?
 - c. A: He had an accident. A car hit him.
 - d. A: He had an accident. ??The car hit him.

While (3c) is perfectly felicitous, (3d) is not. Most theories would predict that the presupposition in (3d) is accommodated globally (i.e., added to the main context), because the result is informative and consistent. So these constraints predict coherence where there isn't any.

These constraints on accommodation also cause difficulties for examples where the presupposition seems to require an explicit linguistic antecedent to be felicitous. Consider, for instance, the contrast between (7) and (8), taken from Beaver (1997):

- (7) ?? I don't know whether the Pope has measles. But every Catholic realizes the Pope has measles.
- (8) I don't know whether the Pope has measles. But for every person, if he's a Catholic and the Pope has measles, then he realizes the Pope has measles.

Van der Sandt's theory predicts that the presupposition that the Pope has measles is globally accommodated in (8). But intuitively, global accommodation shouldn't be possible, because a statement akin to Moore's paradox (1912) would be asserted (*The Pope has measles but I don't know that the Pope has measles*). Van der Sandt's constraints on accommodation don't block this. But even if they did, the resulting account would still allow intermediate accommodation. But this incorrectly predicts that (7) is felicitous. In contrast, the presupposition in (8) is *bound* to this intermediate position. But the difference in acceptability between (7) and (8) demonstrates that an explicit linguistic antecedent is required, and accommodation at the same site as this antecedent is odd. Van der Sandt fails to capture this.

There are certain presupposition triggers, such as *too*, which are like (7) in that their presuppositions require a linguistic antecedent. They pose similar challenges. Consider (9):

- (9) John lived in New York too.

Various people have supposed that (9) presupposes that there is someone other than John who lived in New York. This content is already common knowledge to most speakers. However, (9) is odd in a context where the presupposition or something similar isn't already explicitly introduced into the discourse context. In fact, we think that the received wisdom about the presuppositional content of *too* is wrong: (9) is perfectly felicitous if uttered in a context in which other places John has lived have already been mentioned. The presupposition of *too* is rather that it requires that there be some proposition in the context that bears the rhetorical relation *Parallel* to the content of the sentence in which *too* occurs.³ In any case the presupposition triggered by *too* must be introduced explicitly into the discourse context, and cannot be accommodated.

Another difficulty for this theory is that it on occasion fails to specify the appropriate content of the presupposition itself. Beaver (1997) argues that the following minimal pairs show that the content of the presupposition cannot depend simply on constraints like consistency and informativity:

- (10) If David wrote the article, then the knowledge that no good logician was involved will confound the editors.
- (11) If David wrote the article, then the knowledge that David is a computer program running on a PC will confound the editors.

³For more on *too* and *Parallel*, see Asher (1993) or Asher, Hardt and Busquets (1997).

The presupposition of (10) is (12), whereas for (11) it's (13):

(12) If David wrote the article, then no good logician was involved.

(13) David is a computer program running on a PC.

The factive nominal *knowledge* triggers the presupposition in both (10) and (11). But whether conditional presuppositions arise or not cannot be a structural matter, since the same structure gives rise to a conditional presupposition in (10) but not in (11). So the standard theory of accommodation cannot account for the intuitions about these examples; it doesn't explain the presence of conditional presuppositions in some cases but not in others. Beaver (1997) argues that presupposition satisfaction must also depend on some world knowledge dependent notion of plausibility, but he too doesn't give an explicit or detailed account that would make the correct predictions in (10) and (11).

Van der Sandt's conjecture that presuppositions are anaphors is a compelling idea. But it has shortcomings because accommodation isn't constrained enough. Background knowledge such as information about the domain can block accommodation (Lewis, 1979), and the rhetorical function of the utterance in the text can influence the projection of its presuppositions, as shown with (6). Though Van der Sandt and Beaver would both acknowledge that pragmatic constraints on accommodation are important, these constraints haven't been integrated into the formal theory of presuppositions or made sufficiently precise to have any predictive force. For instance, Van der Sandt's (1992) explicit theory ignores them. Consequently, presuppositions are sometimes accommodated in the wrong DRS to produce the wrong readings, or they are accommodated when they shouldn't be. We will approach these problems by tackling presuppositions in a framework of discourse representation that formalizes the semantics pragmatics interface.

3 The Basic Picture

Our framework of discourse interpretation, SDRT, extends DRT in two fundamental ways. First, discourse contexts are represented as recursive, relational structures (SDRSS), involving DRSS representing the contents of clauses, and discourse relations like *Parallel*, *Narration* and *Background* representing the rhetorical functions of these discourse constituents in the context. Second, these discourse relations can affect the content of the clauses they relate, and hence of the discourse in general.

To build such representations of contexts, SDRT proceeds incrementally, interpreting each bit of new information as yielding a change in the discourse context. Like other dynamic theories of interpretation (e.g., DRT and Groenendijk and Stokhof's (1991) DPL), SDRT must specify how new information can change an existing discourse context—i.e., it must give an account of the context change potential (CCP) of an utterance. SDRT does so as follows. First, one uses the grammar to compositionally build up the DRSS for each clause (cf. Muskens 1995, Fernando 1994, Asher, 1993). The SDRS is then built from these DRSS dynamically. In the simple cases examined here and in Lascarides and Asher (1993), we update clause by clause.⁴

⁴In more complex examples, we may form local SDRSS that are then attached together.

When updating the SDRS built so far with a new DRS, one uses the glue logic to decide where to attach this DRS and to infer one or more rhetorical relations to attach it. This last function of the glue logic was spelled out in Lascarides and Asher (1993) and elsewhere in what we have called DICE (Discourse in Commonsense Entailment), but there the glue logic was restricted to this task. Hence in this paper we will speak more generically and more generally of a glue logic.

The rhetorical relations may trigger modifications to the DRSS that were produced by the grammar. For example, in (14), the glue logic determines that *Narration* connects the clauses. Spatio-temporal constraints on *Narration* are then used to add to the representation of (14) that the boxcar is in Dansville:

- (14) a. John took an engine from Avon to Dansville.
- b. He picked up a boxcar,
- c. and took it to Broxburn.

The rhetorical relations, their semantic effects and the new information are now integrated together with the given SDRS representing the context to form a new SDRS. Thus, SDRS construction proceeds dynamically, and each given SDRS and bit of new information produces a new SDRS, which can serve as a representation of a new discourse context.

Suppose one were to analyse presuppositions from the perspective of this kind of discourse semantics. Then standard binding and accommodation both seem unsatisfactory. First, standard binding is too restrictive: there are many rhetorical relations that bind propositions together—not just identity, but also relations like *Parallel* and *Background*. Binding presuppositions only with identity wouldn't reflect these alternatives. Second, accommodation—where it is viewed as addition alone—is unsatisfactory, because in SDRT the representation of discourse is well-defined or coherent only if each DRS is related to another with a rhetorical relation. There's no scope for just adding a proposition to the context; it has to be related to it as well.

This perspective leads us to loosen van der Sandt's notion of binding. Instead of a presupposition binding with identity to an accessible site in the context, it will bind with a rhetorical relation. For example, in (15), van der Sandt would accommodate the presupposition that there is a King of France, because binding isn't possible.

- (15) The King of France is bald.

We take a different view. (15) means: that there is a King of France is *Background* to his being bald. Much of this paper focuses on spelling out this idea in detail.

Our 'rhetorical' approach to presuppositions has at least three advantages. First, as we mentioned in §1 loosening binding in this way eliminates the need for semantic accommodation altogether, where by semantic accommodation we mean that the presupposition is added to the context, but not related to any element of it. Alternatively, accommodation can be viewed as constrained by much more than informativeness and consistency in this picture, because it is licensed only when the presupposition can be bound to the context with a rhetorical

relation (cf. Lascarides and Oberlander, 1993).⁵ Either way, the accommodation mechanism where information is simply added to the context is removed. This is a very desirable feature of the theory, since this ‘addition’ mechanism is only used to account for presuppositions. We are replacing it with a general procedure of interpretation, which works for assertions too and which is used to model a wide range of linguistic phenomena, particularly semantic ambiguity resolution.

The second advantage of this approach is that it records the influence of pragmatic information on presupposition satisfaction. Pragmatic constraints on computing rhetorical relations, which are provided by the glue logic (Lascarides and Asher, 1993), provide a much more constrained mechanism for dealing with presupposition than informativeness and consistency alone. As a result this theory will overcome some of the problems discussed in §2. Third, it maintains a close relationship between presupposition satisfaction and discourse coherence. The presupposition failure in (3a,b,d) will be predicted by the fact that the presupposition that there is a car cannot be rhetorically bound to the propositions in the context. This is roughly analogous to the discourses given in (16a–c) being incoherent; we can’t attach the presupposition anywhere in the context and get a coherent discourse.

- (3) a. A: Did you hear about John?
 b. B: No, what?
 d. A: He had an accident. ??The car hit him.
- (16) a. A: Did you hear about John? ??There was a car.
 b. B: No, what?
 A: ??There was a car.
 c. A: John had an accident. ??There was a car.

Our account of presupposition will buttress and develop van der Sandt’s view that presupposition satisfaction and discourse coherence are closely related, and we will exploit SDRT’s rich notion of discourse coherence to do this.

In SDRT, assertions are linked to the context by computing a rhetorical relation. We claim that presuppositions are handled this way too. But many linguists have argued that assertions and presuppositions are different. Presuppositions project from embeddings and assertions don’t. And Clark (1977) has intuitively described presupposed information as *given*, and asserted information as *new*.⁶ We need to account for this intuitive difference.

⁵The idea of handling presupposition satisfaction through computing how it is rhetorically connected to the discourse context was used in Lascarides and Oberlander (1993) to process the content of *when-*, *before-* and *after-*clauses. This approach is similar in spirit, but more general. We have replaced their procedural approach with a declarative one; we have integrated semantic and presuppositional information into the construction of an SDRS in a compositional way by using underspecification; and finally, we cover a wider variety of presupposition triggers. Furthermore, Lascarides and Oberlander ignored the way computing a rhetorical relation can lead to further inferences about semantic content, and so they didn’t model the way the content of the presupposition itself is determined by context (e.g., (10) vs (11)).

⁶One could take an even more extreme position, and suggest that presuppositions are extra inferred knowledge, which strictly speaking doesn’t comprise part of the semantic content of the text at all. As Gazdar (1979) and Mercer (1985) have argued, they are implicatures of some kind. We reject this extreme position. Treating presuppositions as part of the semantic content allows one to record how it affects the interpretation of other expressions in the text, such as pronouns (van der Sandt, 1992).

We capture the difference between presuppositions and assertions in three ways. First, the rules for updating discourse ensure that the rhetorical role of presuppositions is in general different from that for asserted information. To capture the intuition that by and large, presuppositions are *given*, the rules for computing rhetorical relations ensure that presuppositions in general bind with *Background* (i.e., the presupposition provides background or sets the stage for the main story line). For example, the representation of (15) will include: (a) the proposition α that x is bald, (b) the proposition β that there is a King of France and that's x ,⁷ and (c) the condition *Background*(α, β) (so the presupposition provides background to the main point, which is that x is bald).

Presuppositions may also bind with *Def-Consequence* in some cases. This relation means that the presupposition is a defeasible consequence of the proposition it attaches to. If the consequence relation is idempotent (a standard assumption), then binding in van der Sandt's sense turns out to be a species of *Def-Consequence*. So binding in van der Sandt's sense will in general correspond in this theory to a case where the presupposition is related by both *Def-Consequence* (because it's defeasibly entailed by the proposition it attaches to) and *Background* (because in our theory presuppositions usually bind this way). Presuppositions can be attached with relations other than *Background*, but in general this arises because the presupposition trigger itself specifies some other rhetorical relation. For example, the presupposition trigger *too* entails *Parallel*, and *because* entails *Explanation*. In contrast to presuppositions, asserted information may be attached to the discourse context with a variety of rhetorical relations other than *Background*—e.g., *Narration*, *Result*, *Explanation*—even in the absence of words like *too* and *because*. Furthermore, assertions don't attach with *Def-Consequence*, unless it's explicitly signalled by a discourse particle.

A second difference between presuppositions and asserted information is that there's an important distinction between their compositional semantics. The grammar produces a representation of presupposed information which is always incomplete or *underspecified*. This is the technical way of ensuring presuppositions are always anaphoric, and must be bound. Because of this, asserted information and presupposed information behave differently at the start of a discourse. Asserted information is simply introduced into the null context. But this isn't possible for presupposed information. It is *always* bound to some antecedent with a rhetorical relation, even if that antecedent is the asserted information in the clause which introduced the presupposition—as in, e.g., (15).

Finally, there will be differences in the preferences for which part of the discourse structure presuppositions vs. assertions attach to. Presuppositions are freer in their attachment possibilities than assertions; those clausal constituents of the assertion that are combined with logical operators have their attachment sites determined by the grammar, whereas presuppositions are always free in principle to attach outside the SDRS constituent in which they were generated. The preferences for attachment predict that presuppositions can project from embeddings in a way that assertions can't.

The underspecification in presuppositions is used to record the influence of pragmatics on presupposition satisfaction. A distinguishing feature of SDRT is that computing a rhetorical connection between propositions in the glue logic can result in those propositions receiving

⁷We ignore here and below the controversial issue as to whether definites introduce a uniqueness requirement (cf. Russell, 1905).

additional content. Update in SDRT is defined so that if this additional content can resolve underspecified conditions that arose in the grammar, then it does. This captures the intuition that when processing discourse, people are expected to and often do fill in gaps in what’s been made linguistically explicit. This feature of SDRT provides a formal model for how pragmatics (which influences the rhetorical relation) affects semantics (i.e., the resolution of underspecification). We will use this feature to record the influence of pragmatics on the semantics of presuppositions.

The grammar typically introduces at least two underspecified elements in a presupposition: one is the rhetorical relation used to bind the presupposed information, and the other is the other term of this relation, which is known as the attachment point. The latter underspecification ensures that the structural position in the discourse context in which the presupposed information is entered is not determined by the syntactic position of the presupposition trigger that generates it, because the attachment point that is part of the representation of the discourse context could be a proposition that was introduced by an earlier sentence (this is explained in more detail in §4.2). This captures the intuition that that presuppositions may project from the embedded contexts in which they are introduced.

Underspecification for presuppositions actually comes in a variety of types, which depend on the presupposition trigger. We will examine some of these in §4.1. In all cases, however, in processing a presupposition one must resolve the underspecified elements with respect to the discourse context. So presupposition projection, binding and accommodation all occur as a byproduct of the SDRT update procedure. They occur by computing how the presupposition rhetorically connects to the context. Since this update procedure is determined by the pragmatic information specified in the glue logic, presuppositions are influenced by pragmatics as well as compositional semantics.⁸

4 Presuppositions in More Detail

Having given the basic picture, we will now turn to the details of the theory. First, we will give an overview of the SDRT compositional semantics of presupposition triggers. Processing the presupposition relative to the context amounts to resolving its underspecified elements, and this is handled via the SDRT update function. We will illustrate how this works, by analysing some simple sentences (i.e., sentences that don’t contain logical operators such as *if* or *not*). We will then examine the interpretation of presuppositions in more complex sentences (e.g., (10) and (11)), and compare our treatment to that of other dynamic accounts, such as Beaver’s (1996) and van der Sandt’s (1992). Using SDRT makes comparisons and reformulations of van der Sandt’s DRT-based theory of presuppositions relatively straightforward. We will demonstrate that at least some of the problems concerning conditional presuppositions can be overcome, because of the role of pragmatics in SDRT. Finally, we will examine presuppositions in multi-sentence discourse, and then return to some cases that one might think are difficult to handle in a proposal where semantic accommodation is excluded.

⁸Indeed the principal advantage of using SDRT for investigating presupposition over other theories of discourse structure (Hobbs et al. 1993, Grosz and Sidner 1986, Polanyi 1985, and others) is that SDRT has already explored extensively the interaction between compositional semantic information extracted from the grammar and pragmatic information.

4.1 Semantics of Clauses

We first examine some important aspects of the compositional semantics of the clauses. Following Reyle (1993) and Asher and Fernando (1997), the grammar utilises underspecification to represent semantic ambiguity. The grammar produces a set of labeled pieces of information and a set of constraints about how these labels can combine. Semantic ambiguity occurs when the constraints underdetermine the combination of these labels. Since the combination of labels is underspecified, so is the representation of the meaning of the clause. Resolving the relative scopes of the labels is then done by updating the discourse context with this (underspecified) meaning. Discourse update in SDRT is specifically designed to resolve underspecified elements when it can.

To avoid clutter in notation, semantic information that must be grouped together according to syntax receives just one label. Moreover, as in HPSG, (Pollard and Sag, 1994), these labels may appear in conditions. A special case of this labelling procedure is the following: if the grammar produces the DRS K_π for some clause, then the SDRS for the discourse will include (17) where the label π is a discourse referent, which we call a ‘speech act discourse referent’ since it labels a DRS (we will always assume a notational convention that π labels K_π):

$$(17) \quad \boxed{\begin{array}{c} \pi \\ \pi : K_\pi \end{array}}$$

Updating the discourse with information such as that in (17) will involve computing a rhetorical relation R between π and some accessible speech act discourse referent π' in the discourse context. Computing this rhetorical relation may include resolving underspecified conditions that are in K_π , as we’ll see in §4.2. Note that R ’s arguments are the labels, and not the DRSS themselves. We explain why in §4.2.

A presupposition that’s introduced by a clause can receive different scope from the assertion. For example, in (1), the presupposed information has wider scope than the conditional, but the asserted information does not.⁹ In §2, we showed that the relative scopes of the asserted and presupposed content of the clause isn’t syntactically determined; and so, in line with the general strategy for representing ambiguity, the grammar groups asserted information under one label, and the presupposed information under another. Consequently, the grammar produces an SDRS like that in (17) for the asserted content of a clause, and it will produce another SDRS with speech act discourse referent π' , and condition $\pi' : K_{\pi'}$ for the presupposed content of this clause, where $K_{\pi'}$ will be the DRS discourse constituent that represents the presupposition (e.g., *Jack has a son* for (1)). The grammar doesn’t determine the scope of π and π' . Rather, all of these SDRSS must be attached to the representation of the discourse achieved so far through the SDRT discourse update procedure: that is, they must be attached to the context with a rhetorical relation. Rhetorical relations produce hierarchical structures for the labels π , π' etc, since some rhetorical relations are subordinating and some are coordinating. So this discourse structure will resolve the relative scope of the asserted and presupposed content.

⁹Since our main concern is modelling presuppositions, we will ignore problems of quantifier scope and concentrate only on the different possible ‘scopes’ of presupposed and asserted information.


As well as separating asserted content and presupposed content in this way, and providing them with different labels, the grammar distinguishes between presuppositions and assertions in one other important respect. Presuppositions are explicitly encoded as anaphoric, since the grammar invokes for each presupposition an underspecified attachment point and underspecified rhetorical relation. These are given respectively by $u = ?$ and $R = ?$ in the canonical representation of presuppositions, given in (18):

$$(18) \quad \boxed{\begin{array}{l} \pi', u \\ \pi' : K'_\pi \\ R(u, \pi') \\ R = ? \\ u = ? \end{array}}$$

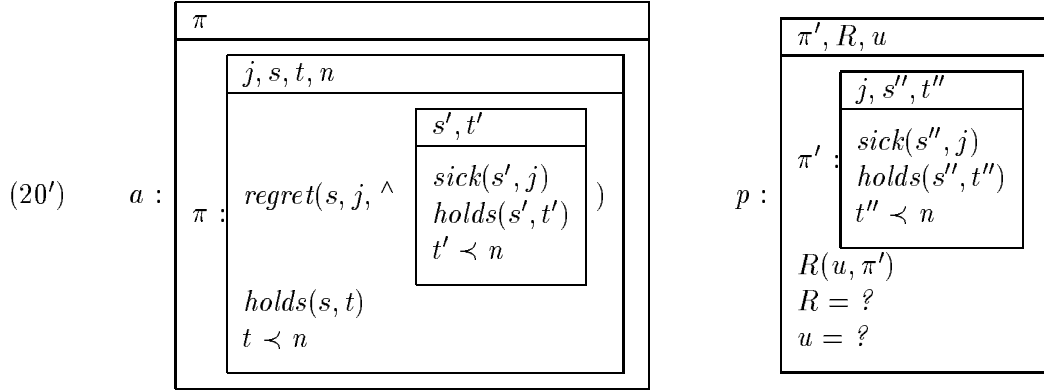
This distinguishes this treatment of presuppositions from van der Sandt’s, because this underspecification means that the presuppositions must *always* be rhetorically bound to the context, rather than added.

The Projection Problem is challenging, because the scope of presupposed information is contextually rather than grammatically determined. Our semantic representation of presuppositions allows for this through the use of the underspecified attachment site $u = ?$. The scope of a presupposition is determined by resolving $u = ?$. Since this is done via the discourse update procedure, and hence via the glue logic, pragmatics will influence the scope of presuppositions. Moreover, for presupposition triggers that are introduced in an embedded context such as a conditional, SDRT allows the presupposition to access attachment points that are outside the scope of the embedding. If discourse update resolves the antecedent attachment point u to one of these, then the presupposition has essentially projected out from its embedded context. This contrasts with the assertions in the conditional, which can’t project in this way.

To illustrate the ideas, let’s examine the compositional semantics of some example clauses that contain presupposition triggers. Presuppositions of *abstract type*—i.e., propositional or factual presuppositions—are triggered by expressions such as *wh*-questions, focus, comparatives, *stop/start* verbs, *manage* and *succeed* verbs, factive verbs like *know*, *regret*, *realize*, and presuppositional adverbs like *again*. Let’s consider, for instance, a factive verb like *regret*. The grammar will produce two SDRSs for a sentence containing *regret* as shown in (19): a is the asserted information, and p is the presupposed information; [1] is filled in with the subject NP and [2] with the sentential complement of *regret*. So the grammar produces the representation (20’) for sentence (20). a and p below each label bits of structure that are produced by the grammar, but they’re *not* discourse referents themselves. They merely provide ways of talking about the whole structures, and therefore act as guides for putting the SDRSs together; they will disappear after SDRS update. This will involve binding π , π' etc to available referents with rhetorical relations. Note also that the DRS representing *John is sick* forms part of the representation of the asserted content, and also part of the representation of the presupposition. We use the DRS in a rather than the label because the intensional operator works over a DRS and not a term.




(20) John regretted that he was sick.



The SDRS p contains the underspecified conditions $R = ?$ and $u = ?$, which must be resolved by updating the context with p ; i.e., by computing a rhetorical relation between π' and an antecedent site in the context. Resolving the attachment site u will determine the scope of the presupposition.

As we've mentioned before, some cue phrases are presupposition triggers which actually specify the rhetorical relation (e.g., *too* specifies the relation to be *Parallel*). The presupposed information triggered by *too* and *also* is represented in (21), where [1] is to be filled in by a label for the SDRS over which *too* (or *also*) is assigned scope by the grammar:




Such cue phrases presuppose that an antecedent of an appropriate sort exists. This is represented by the anaphoric condition $u = ?$. Other cue phrases may introduce different discourse relations—e.g., *because* introduces *Explanation*; *although* introduces *Contrast*. All of these will generate the same sort of conditions as (21). Such conditions impose constraints on what u can be, because each rhetorical relation specifies constraints on the relationship between the semantic content of the propositions it connects. These constraints are described in §4.2.

Let's consider now the presuppositions of definites. In line with Chierchia (1995) and Asher and Lascarides (in press), we assume that definites introduce an underspecified bridging relation, which connects the object denoted by the definite to some antecedent in the context. In Asher and Lascarides (in press), we demonstrated that this bridging relation could be

computed as a byproduct of SDRT Update, which is how we also treat presuppositions.¹⁰


Definites introduce not only underspecified ‘bridging’ relations at the ‘micro-structure’ level to connect the denoted objects together, but in line with other presupposition triggers, they also introduce underspecified ‘rhetorical’ relations at the ‘macro-structure’ level of the from $R(u, \pi)$, $R = ?$ and $u = ?$, which connect the propositions together:



The definite NP furnishes argument [1], while the VP contributes to argument [2] (depending on whether the definite occurs in subject or object position).

As an illustrative example, sentence (23) has the compositional semantics given in (23'):

(23) The man walked



The presupposition must be bound to the context via a rhetorical relation. In addition, the


¹⁰ Asher and Lascarides (in press), however, ignored spatio-temporal conditions on when the bridging relation holds. We rectify this here, by subsuming the bridging within our general treatment of presuppositions. We'll demonstrate that resolving the rhetorical connection between the presupposed content and the discourse context introduces spatio-temporal constraints on that content, which in turn impose the spatio-temporal conditions on the bridging relation B in the required way.

man denoted by the definite must be (bridging) related to an antecedent object (so (23) couldn't be uttered in a null context).

Our treatment of definites extends to quantified NPs, which also convey presuppositions (cf. Von Stechow 1994). The restrictor includes underspecified bridging conditions, which require the objects denoted by the NP to be related to antecedent objects. This is intended to capture the bridging that occurs in discourses such as (24): the strikers mentioned in (24b) played in the football match mentioned in (24a):

- (24) a. I went to the Scotland vs. Latvia football match yesterday.
 b. Every striker attempted to score a goal.

We reflect this in the semantics of quantified NPs by invoking underspecified conditions of the type found in definites, which are placed in the restrictor of the DRS duplex condition. For example, the compositional semantics of *every* is specified in (25), where [1] is derived in the grammar from the NP that contains the determiner *every*, and [2] is derived from whatever the NP combines with to form a clause:



In (24) for example, [1](*x*) is replaced with *striker(x)*; [2](*x*) is replaced with the VP *attempted to score a goal*, and *B*, *y* and *u* will be computed through SDRT Update, so that the quantified NP connects with (24a) in the right way. That is, *B* will be resolved to *play*, *y* identified with a discourse referent denoting the football match, *u* identified with the label for the DRS which represents (24a), and *R* resolved to *Background* (as we will see shortly). Note that unlike definites, the presupposition associated with *every* is restricted in scope to being in the restrictor; it isn't free to take wider (or in DRT, 'higher') scope. This restriction ensures that *every striker V-ed* doesn't entail that there are any strikers. Of course, these observations hold provided that the discourse referents within the constituent labelled by π get the appropriate quantificational force from \Rightarrow . This is assured if *R* is a veridical relation (cf. §4.3), because then any embedding of the right hand side of the conditional must perforce be a proper embedding of the π constituent.

These examples illustrate the variety of underspecifications that presupposition triggers generate. The update task involves resolving the underspecified conditions where possible. This includes scope resolution. So using underspecification to represent the compositional semantics of presuppositions allows us to assimilate presupposition projection to a more general problem—that of assigning a scope to various informational units.

We now turn to the problem of defining discourse update, and the way discourse update resolves underspecification.

4.2 From Clauses to Discourse

The SDRSS produced by the grammar must be dynamically integrated together, into a representation of the semantics of the discourse. In SDRT, one updates the discourse context τ that’s been built so far with some new information β , by attaching β to α with a rhetorical relation, where α is an available attachment site (i.e., speech act discourse referent) in τ .

But which α ’s are available attachment sites? Rhetorical relations provide a richer discourse structure than DRT alone gives: *Elaboration*, \Downarrow ($\alpha \Downarrow \beta$ can be glossed as “ α is a topic for β ”) and *Explanation* are subordinating relations; the rest are not. Therefore, the notion of accessibility in SDRT (which is termed availability to distinguish it from DRT) is more constrained than in DRT. Available labels in an SDRS are those on its right frontier: that is, the previously attached label, and ones which elaborate or explain it, or are a topic to it. So, for example, new information could attach to π' in (26) but not to π :

$$(26) \quad \boxed{\begin{array}{l} \pi, \pi' \\ \hline \pi : K_\pi, \pi' : K_{\pi'} \\ \text{Narration}(\pi, \pi') \end{array}}$$

And hence the antecedents to anaphora must be picked up from $K_{\pi'}$, but not from K_π . In general, this offers a much smaller set of possibilities than DRT itself would offer. This is one way in which interpreting presuppositions will be more constrained in SDRT than in DRT.

To compute which rhetorical relation to use, it is essential to use a nonmonotonic, ‘glue’ logic, as we have argued in many places (e.g. Lascarides and Asher 1993). It is also essential to make use of nonmonotonicity to place constraints on *where* to attach new information, especially when this information to be attached is presupposed. The glue logic is expressed in a quantifier-free language augmented with a weak conditional connective: $A > B$ means “If A , then normally B ”. This glue logic for constructing logical forms for discourse is quite simple, even though the content of the discourse may be very complex. To build logical forms, the glue logic must have access to the contents of the discourse, but to keep the logic simple it has only limited access to the full contents of the propositions—namely the SDRSS. More specifically, we suppose that the glue logic has access to the form of the SDRSS, but not what they denote (Asher and Fernando, 1997). To this end, we suppose that there is an information transfer function μ from the language of information content to the glue language, which takes conditions inside a given SDRS and turns them into predicates on that SDRS’s label. So $\mu(K_\alpha)(\alpha)$ is a set of formulae of the glue language of the form $\phi(\alpha)$, where

ϕ is a predicate on the propositional variable α , and it corresponds to an SDRS condition in K_α .

The glue language uses various nonlogical predicates in describing rules which constrain where new information should attach, and what rhetorical relation should be used to attach it. We exploit the labels of SDRSs in the glue language, as arguments to our nonlogical predicates (Asher, 1996). One of these is $\langle \tau, \alpha, \beta \rangle$, which can be glossed as “the constituent labelled by β is to be attached to the constituent labelled by α with a rhetorical relation, where α is an available label in the SDRS labeled τ that’s built so far”. There are also relation symbols for the discourse relations, whose arguments are the labels (e.g., *Explanation*(α, β)); predicates on labels, which allow us to specify whether the main eventuality described in the SDRS with that label is an event or a state (e.g., *event*(α) and *state*(α)); and predicates on the labels, which describe conditions within the SDRSs (e.g., *overlap*(e, e')(α)). For example, the **Narration** axiom of our glue logic below states: If the constituent labelled by β is to be attached to the constituent labelled by α with a rhetorical relation, and those constituents describe events, then normally, the rhetorical relation is *Narration*. On the other hand, if β is stative, then by **Background**, the rhetorical relation is normally *Background*:

- **Narration:** $(\langle \tau, \alpha, \beta \rangle \wedge \textit{event}(\alpha) \wedge \textit{event}(\beta)) > \textit{Narration}(\alpha, \beta)$
- **Background:** $(\langle \tau, \alpha, \beta \rangle \wedge \textit{state}(\beta)) > \textit{Background}(\alpha, \beta)$

For example, **Background** ensures that *Background* relates the contents of the first and second sentences in (27):

(27) Jack has a son. He is bald.

The language also includes a predicate \downarrow and a function symbol *Update*. $\textit{Update}(\tau, \alpha, \beta)$ returns the SDRS constructed by the glue logic given an input SDRS τ , an attachment point α in τ , and a new labelled SDRS β . $\downarrow \textit{Update}(\tau, \alpha, \beta)$ means that updating τ with β by attaching it to α produces a well-defined SDRS (i.e., there are no unresolved conditions $x = ?$, and every constituent is rhetorically attached to another). We will define SDRS Update precisely in a minute and use this to specify constraints on which site α new material β should attach to. Once again, it is important to stress that while the glue logic has access to the structure of SDRSs—e.g., the conditions within them—and indeed reasons about SDRSs, it does not have access to their denotations. The glue logic reasons about logical forms of discourses, not (except in the very limited way we have mentioned) about their content.

Because the majority of the glue logic axioms feature the connective $>$, and $>$ doesn’t support a modus ponens like rule, computing rhetorical relations is nonmonotonic. The nonmonotonic notion of validity \approx that underlies the SDRT glue logic has several desirable properties, which we have discussed in detail elsewhere (Lascarides and Asher, 1993). There are three that are relevant for our purposes. First, \approx validates Defeasible Modus Ponens (DMP): when the default laws whose antecedents are verified all have consequents that are consistent with the monotonic information and with each other, then all the consequents are nonmonotonically inferred. Second, \approx validates Specificity: when conflicting default rules apply, the consequent

of the most specific default rule (if there is one) is inferred. And finally, \approx is robust in that if $\Gamma \approx \phi$, then $\Gamma \cup \psi \approx \phi$, where ψ is logically independent information.

The glue logic also contains (monotonic) axioms concerning the semantic effects of rhetorical relations. For example, **Axiom on Narration** stipulates that any items that are connected with *Narration* must describe eventualities where the first precedes the second ($e_\alpha \prec e_\beta$ means α 's main eventuality precedes β 's, and it's a gloss for a formula of the propositional glue language). The **Axiom on Background** stipulates that the eventualities overlap. And the **Spatial Consequence of Narration** is derived from certain commonplace assumptions about eventualities (Asher *et al.*, 1996): if *Narration*(α, β) holds and α and β share an actor x then the location of x is the same at the end of the event e_α as it is at the beginning of the event e_β :¹¹

- **Axiom on Narration:** $Narration(\alpha, \beta) \rightarrow e_\alpha \prec e_\beta$
- **Axiom on Background:** $Background(\alpha, \beta) \rightarrow overlap(e_\alpha, e_\beta)$
- **Spatial Consequence of Narration:**
 $(Narration(\alpha, \beta) \wedge actor(x, \alpha) \wedge actor(x, \beta)) \rightarrow$
 $loc(x, source(e_\beta)) = loc(x, goal(e_\alpha))$

These rhetorical relations and their inferred semantic effects are used to define the appropriate update function in SDRT.

As well as axioms for determining a rhetorical relation, we also need to encode constraints on where to attach, because in general more than one attachment site is available in the SDRS τ that's been built so far (e.g., if the previous constituent α_2 was attached to α_1 with *Explanation*, then both α_1 and α_2 are available, and new information β could potentially bind to either one of these). We consider this issue in the particular case when the information to be attached is a presupposition. As many authors have noted, there is the intuition that presupposed information conventionally portrays given information in some sense (Clark, 1977): assigning the presupposed information wide scope in SDRT may very well place it in an earlier part of the discourse context than the one containing the presupposition trigger. To reflect this, we will encode a default version of van der Sandt's preference for global accommodation (cf. Mercer, 1987). This is an important distinction. In van der Sandt's account, the preference for global accommodation always wins, unless the constraints of informativeness and consistency block it or the information can be bound at some lower site. In our theory, the exceptions to the default preference aren't enumerable. Any conflicting monotonic rule or more specific default will override it. We'll see examples of such rules shortly, and also in §5.2. Furthermore, the default preference can be encoded declaratively in SDRT, in contrast to Van der Sandt where the way presuppositions are processed isn't part of the declarative definition of DRT Update.

The declarative rule, which captures a default preference for global accommodation, is given below. First some notation: If β labels asserted content, then $\partial(\beta)$ is the label for the

¹¹This spatial constraint applies only to *Narration* because, as we have noted elsewhere (Asher and Lascarides, 1995), it does not apply for all discourse relations. For example *Contrast* and *Parallel* have different spatial constraints, as given in Asher and Lascarides (1995, in press).

corresponding presupposed content. So **Prefer Global Attachment** says: if α is the most superordinate node in τ to which the presupposition $\partial(\beta)$ can attach, then normally, it attaches there:¹²

- **Prefer Global Attachment**

If α is the most superordinate node in τ , then normally: $\langle \tau, \alpha, \partial(\beta) \rangle$

This default will ensure that global attachment occurs in (1) and (11). However, as we just mentioned, making this rule default allows other constraints on SDRT update to override it. One particular rule that can do this is a constraint called **Maximize Discourse Coherence**, which intuitively captures the idea that people tend to interpret text in the way that makes ‘most sense’. We have argued elsewhere that such a rule is needed to explain data where ambiguity resolution and bridging inferences are influenced by a complex interaction between semantics and pragmatics (e.g., Lascarides, Copestake and Briscoe 1996, Asher and Lascarides in press). The rule we give below is a generalization of the rules in those papers. We assume the following notational convention: $\beta \rightsquigarrow_* \beta_1$ means that K_{β_1} is a DRS which represents one way of resolving the underspecification in K_{β} .

As background to our rule, we assume that in some contexts, attaching with one kind of rhetorical relation produces a ‘closer connection’ or ‘better coherence’ than attaching with another kind (see Asher and Lascarides, in press). To express this, we introduce a partial order on rhetorical relations, relative to the content of the context: *Explanation* \succ_{τ} *Background* means that it would be preferable to interpret the new information that is to be attached as an explanation in τ , rather than as background information. Both alternatives may be coherent, but one is better than the other, and this is partly because of the content of τ . But it won’t do just to look at the strength of the relation which attaches new information, perhaps resolved or not, to some site in τ . We have to compare the strength of all the discourse relations that are affected by the update. This means that we have to compare various updates—viz. those using different attachment points or different resolutions of underspecified elements in the information β to be attached to the discourse context. So we will introduce a predicate that says that a particular update is (coherence) maximal with respect to the new information and the context, which, for a given SDRS τ and new information β , we write as τ, β -maximal. The semantics of this notion is given in terms of the above partial order on relations, and comparing pairs of relations in any two updates under a partial, structural isomorphism ϑ whose kernel on all updates is τ .¹³ We say that an update, $Update(\tau, \alpha_1, \beta_1)$, is τ, β -maximal iff $\beta \rightsquigarrow_* \beta_1$ and for any update $Update(\tau, \alpha_2, \beta_2)$ such that $\beta \rightsquigarrow_* \beta_2$: (i) there is a relation R in $Update(\tau, \alpha_1, \beta_1)$ and $\vartheta(R)$ in $Update(\tau, \alpha_2, \beta_2)$ such that $R \succ_{\tau} \vartheta(R)$ and (ii) for every relation R in $Update(\tau, \alpha_1, \beta_1)$ and $\vartheta(R)$ in $Update(\tau, \alpha_2, \beta_2)$, $R \succeq_{\tau} \vartheta(R)$. Since \rightsquigarrow_* resolves all underspecifications present in the SDRS, maximal updates are necessarily well-defined.

The rule below then captures the following: the underspecified elements in K_{β} are normally resolved so as to maximize discourse coherence, regardless of which attachment site is the most superordinate one:

¹²For simplicity, we have written this rule in words, although it can be encoded using the modal connective \triangleright .

¹³This kernel ensures that only the relations in the context change, and not the structure.

• **Maximize Discourse Coherence:**

- If (a) $\beta \rightsquigarrow_* \beta_1$; and
 (b) $Update(\tau, \alpha, \beta_1)$ is τ, β -maximal; and
 (b) α' is the most superordinate node in τ

Then normally:

- (c) $\langle \tau, \alpha, \beta_1 \rangle$

First, note that this rule has a more specific antecedent than **Prefer Global Attachment**. So it will override that rule, if it conflicts with it. But what circumstances will produce conflict between these two rules? Suppose that β labels a presupposition, and therefore there is a condition $u = ?$, which encapsulates that the attachment point to which β should be rhetorically linked is underdetermined. Suppose furthermore that there is a (local) attachment site α_l which yields a discourse relation R_1 between α_l and β , whereas the global attachment site α_g results in the rhetorical relation R_2 for linking β to α_g . Suppose furthermore that R_1 is maximal in the partial order. Then **Maximize Discourse Coherence** favours resolving the underspecification in β so that $u = ?$ resolves to $u = \alpha_l$, and $R = ?$ resolves to $R = R_1$. This means that the attachment point α_l is preferred over the global site α_g according to this rule. However, **Prefer Global Attachment** applies where α_g replaces α . And so **Prefer Global Attachment** yields a conflicting default preference, that the presupposition β attach to α_g rather than α_l . Given that **Maximize Discourse Coherence** is more specific, the local attachment will be chosen over the global one in these circumstances.

This interaction between **Prefer Global Attachment** and **Maximize Discourse Coherence** will explain why the presupposition that the problem's been solved in (6) is bound locally, in spite of the default preference for 'presupposition projection'.

- (6) a. A: The Problem Solving Group is given a problem each day, and the group leader Mary has to assign it to someone in her group. John is the best problem solver. But when he solves a problem, he always boasts about it. This annoys Mary, and so if she thinks that the day's problem is an unsolved one, she gives it to him, to test him. Otherwise, she gives it to someone else.
- b. B: John's being very quiet just now. Did she give him today's problem?
- c. A: Well, I'm not sure she did. Either John didn't solve the problem or else Mary realized that the problem's been solved.

We'll give a detailed analysis of this in §5.3. We'll also see how **Maximize Discourse Coherence** predicts local accommodation for the conditional sentence (10) in §5.2.

- (10) If David wrote the article, then the knowledge that no good logician was involved will confound the editors.

Arguably, the most important feature of SDRT update is that it models the way the content of a discourse is more than what's explicitly stated linguistically. It does this as follows: if the DICE axioms yield a nonmonotonic conclusion that the rhetorical relation between α

and β is R , and information that’s necessary for the coherence of R isn’t already in K_α or K_β (e.g., $Background(\alpha, \beta)$ is nonmonotonically inferred, but $overlap(e_\alpha, e_\beta)$ is not in K_α or in K_β), then this content is added to the constituents in a constrained manner through the SDRS update process. Hence the result of an updating procedure can be not only the addition of a rhetorically connected constituent, but also the addition of semantic content to that constituent, that the grammar didn’t provide.

We’re now in a position to define the update procedure in SDRT. Update in DRT (glossed as $Update_{DRT}$) is used to define update in SDRT (glossed as $Update_{SDRT}$). $Update_{DRT}$ uses set union:¹⁴

- **DRT’s Update Function:**

$$Update_{DRT}(K_1, K_2) = [U_{K_1} \cup U_{K_2}][C_{K_1} \cup C_{K_2}]$$

When the discourse context is “empty” (i.e. when we are just starting a discourse), $Update_{SDRT}$ is just $Update_{DRT}$. But in general, $Update_{SDRT}$ is much more complex. In particular, when the discourse context τ is a non-empty SDRS, $Update_{SDRT}$ requires an available attachment point α in τ to be identified, for attaching the new information β to with a rhetorical relation. In general, there is some indeterminacy in this choice of α in spite of the default preferences given in DICE, which constrain but may not fully determine which available site to attach new information to. So to specify the update of old information with new, we need to specify three arguments: K_τ for the old information, K_α (which is on the right frontier of K_τ) for the attachment point, and K_β for the new information. We write this as $Update_{SDRT}(K_\tau, K_\alpha, K_\beta)$.

Moreover, we don’t perform set union on K_α and K_β . Rather, they’re brought together by computing a rhetorical relation between α and β in DICE.

Were we to ignore the effects of discourse structure on semantic content, and in particular on the resolution of underspecified conditions, we could leave the matter here: a non-empty discourse context with an attachment point α , $Update_{SDRT}(K_\tau, K_\alpha, K_\beta)$ would be a new SDRS $K_{\tau'}$, which is just like K_τ , except that it also includes: (a) the new information K_β ; and (b) an attachment of β to α with a rhetorical relation R that’s computed via DICE.

However, it is a central point of our theory that the choice of a rhetorical relation does affect semantic content, and the resolution of underspecified conditions. The update of a discourse context should therefore ensure that the content of K_τ , K_α and K_β are modified, so that the coherence constraints on R are met (where $R(\alpha, \beta)$ is inferred via DICE), and underspecified conditions are resolved where appropriate. In fact, SDRT Update is defined so that additional semantic content φ , that’s inferrable in DICE from the discourse context and K_β , is added to K_β in the update. In addition, φ can be used to resolve underspecified conditions in K_β , K_α , and any $K_{\alpha'}$ such that α' is attached to α with a rhetorical relation.

A final complexity comes with the need to distinguish several cases. First, we need to specify how updates of a null context with new asserted information proceed (clause 1 below). This case is distinct from the others, because the information is added to the null context, rather

¹⁴Note that this definition holds only for DRSS without conditions of the form $x = ?$. We gloss over the complications that occur when we countenance such conditions. They aren’t relevant for our purposes, since we resolve conditions such as these in different ways from those in DRT.

than linked to it with a rhetorical relation. We then need to specify updates of a non-null context with asserted information (clause 2), and with presupposed information (clause 3 below). These are different, because of the underspecified rhetorical conditions in presuppositions, which are absent from asserted information.

We now define $Update_{\text{SDRT}}$ formally. As we have just seen, there are three cases: one when the discourse context is empty, one when asserted information like (17) is attached to a non-empty discourse context, and one when presupposed information (with underspecified rhetorical conditions) is attached to a non-empty discourse context. Let \sim be the (nonmonotonic) proof theoretic consequence relation for the glue logic. Let $Pred_{\pi}$ be the label of the SDRS constituent in which π is declared, or (equivalently) in which a condition of the form $\pi : K$ occurs. And let $\alpha[\beta/\gamma]$ be the result of replacing γ in α with β . Then $Update_{\text{SDRT}}$ is defined as follows:

• **SDRT's Update Function:**

1. $Update_{\text{SDRT}}(\emptyset, \emptyset, [\beta][\beta : K_{\beta}]) = Update_{\text{DRT}}(\emptyset, [\beta][\beta : K_{\beta}])$
2. For $K_{\tau} \neq \emptyset$, $Update_{\text{SDRT}}(K_{\tau}, K_{\alpha}, [\beta][\beta : K_{\beta}])$ is the SDRS $K_{\tau'}$ such that:
 - (a) $(\langle \tau, \alpha, \beta \rangle, \mu(K_{\alpha})(\alpha), \mu(K_{\beta})(\beta)) \sim (R(\alpha, \beta) \wedge \varphi)$; and
 - (b) $K_{\tau'} = K_{\tau}[K^{+}/K_{Pred(\alpha)}]$, where:
 - (c) $K^{+} = Update_{\text{DRT}}(K'_{Pred(\alpha)}, [\beta][\beta : K_{\beta}(\varphi), R(\alpha, \beta)])$, where $K_{\beta}(\varphi) = K_{\beta}$ together with those conditions specified by φ (where φ is \sim -provable) that are needed to satisfy the coherence constraints on R ; where these conditions φ are also used in $K_{\beta}(\varphi)$ to resolve underspecified conditions in K_{β} (e.g., $x = ?$); and
 - (d) $K'_{Pred(\alpha)}$ is just $K_{Pred(\alpha)}$ except that it may contain specifications of underspecified conditions in $K_{\alpha'}$, such that $\alpha' = \alpha$, or for some rhetorical relation R , $R(\alpha', \alpha)$ or $R(\alpha, \alpha')$ are conditions in $K_{Pred(\alpha)}$, if those specifications are \sim provable from $(\mu(K_{\alpha})(\alpha), \mu(K_{\beta})(\beta))$ and $R(\alpha, \beta)$.
3. For $K_{\tau} \neq \emptyset$, $Update_{\text{SDRT}}(K_{\tau}, K_{\alpha}, [\beta][\beta : K_{\beta}, R_0 = ?, u = ?, \dots])$ is the SDRS $K_{\tau'}$ such that:
 - (a) $(\langle \tau, \alpha, \beta \rangle, \mu(K_{\alpha})(\alpha), \mu(K_{\beta})(\beta)) \sim (R(\alpha, \beta) \wedge \varphi)$; and
 - (b) $K_{\tau'} = K_{\tau}[K^{+}/K_{Pred(\alpha)}]$, where:
 - (c) $K^{+} = Update_{\text{DRT}}(K'_{Pred(\alpha)}, [\beta][\beta : K_{\beta}(\varphi), R_0 = ?, u = ?, \dots])[R_0 = R/R_0 = ?][u = \alpha/u = ?]$, where $K_{\beta}(\varphi)$ is specified as above, and
 - (d) $K'_{Pred(\alpha')}$ is specified as above.

In words, clause 1 deals with the case when we start a new discourse. It applies only to the asserted information of a clause, for the new information doesn't contain underspecified rhetorical conditions $R = ?$ (in contrast to clause 3). It applies only to asserted content, because clause 1 simply adds this to the SDRS for the discourse. Doing this with a presupposition would produce an SDRS which contains the unresolved underspecified conditions $R = ?$ and $u = ?$, resulting in a representation that's not well-defined. Consequently, at the start of a discourse, one must attach the asserted information *before* one updates this with

the presupposed information (using clause 3). However, when continuing a discourse rather than starting it, the relative order for attaching the asserted or the presupposed content of a sentence isn't specified. We'll say more on this shortly.

Clauses 2 and 3 deal with the case when we continue with a discourse (because K_τ is non-empty). Clause 2 deals with asserted information (because the SDRS β doesn't contain conditions of the form $R = ?$ and $u = ?$), and clause 3 deals with presupposed information (because these conditions are present in β). Let's look at clause 2 first. Clause 2 demands that you rhetorically connect β to α . Clause (2a) ensures that the rhetorical relation R that's chosen is the one that results from the DICE axioms. As can be seen from clauses (2b) and (2c), $R(\alpha, \beta)$ gets added to the update, specifically to the constituent that introduced the attachment point α as a speech act discourse referent. Furthermore, φ in clause (2a) stands for information that follows from DICE when reasoning about how β attaches to α . Thus φ will include information additional to that given by the grammar, that follows from R 's coherence constraints (e.g., if R is *Background*, φ includes *overlap*(e_α, e_β)). Clauses (2b) and (2c) guarantee that this additional semantic content that's inferred via DICE also gets added, this time to the constituent K_β that was derived via the grammar. Moreover, (2c) stipulates that φ replaces underspecified conditions in K_β with specified ones (in $K_\beta(\varphi)$), if it can (we'll see many examples where this happens in §5). Clauses (2c) and (2d) also ensure that information φ that's inferred in DICE resolves as many underspecified conditions as it can in $K_{Pred(\alpha)}$. More precisely, underspecified conditions in the attachment K_α and things that are rhetorically linked to it can be resolved by φ .

Clauses (3a–d) stipulate that updating a discourse with a presupposition is just like doing it with assertions (i.e., clauses (2a–d)): that is, a condition $R(\alpha, \beta)$ is added; K_β gets modified with additional information φ ; and K_β , K_α and $K_{\alpha'}$ (where α' is attached to α) can have underspecified conditions resolved by φ . But there's one more thing: the underspecified conditions $R_0 = ?$ and $u = ?$ are replaced with $R_0 = R$ and $u = \alpha$. In other words, the rhetorical anaphoric elements in the presupposition reflect the rhetorical connection and attachment point that's given by DICE.

Note that $Update_{SDRT}$ has been defined for attaching only one packet of information to the discourse structure. But the grammar produces more than one packet of information for each clause: there's one for the asserted content, and one for the presupposition. So defining how to update a discourse with the content of a clause involves refining the definition of SDRT update, to allow (perhaps) several attachments and their consequent modifications to the content. For the purposes of this paper, we'll assume that updating the discourse with a clause is defined with respect to a *set* of attachment points—one for the asserted information and one for each distinct presupposition. This more complex function is simply the result of sequentially carrying out the procedure of $Update_{SDRT}$ as defined above, for the asserted information and each presupposition according to some sequence. In contrast to van der Sandt, we remain agnostic as to whether updating with presupposed information or asserted information is done first (except for the first clause in the discourse). The result of any sequence of updates will be viewed as an admissible integration of the new information in the discourse.

Overall, the definition of SDRT update guarantees that the way a presupposition is interpreted in discourse is dependent on a complex mixture of compositional semantic and pragmatic

information. This is because the DICE axioms that are used to infer R (in clause 3) are based on linguistic and non-linguistic knowledge. Moreover, the interpretation of presuppositions is constrained by a richer notion of accessibility than that given by DRT, since presuppositions can attach to what’s SDRT-available, which in general is a smaller set than what’s DRT-accessible. We’ll demonstrate in our analysis of examples that these two features overcome some of the problems with traditional dynamic approaches to presuppositions. Furthermore, the above axioms in the glue logic, which constrain which attachment site to use, will ensure that updating the context with a presupposition is in general different from doing it with an assertion, in the manner we described in §3. In particular, these axioms will determine which attachment sites presuppositions should attach to, and normally, they will be higher in the structure than where the corresponding assertion attaches. In other words, presuppositions will tend to have wider scope than assertions.

The definition of Update models the resolution of ambiguity, because it constrains the updated SDRS to be one where inferences in DICE resolve underspecified conditions whenever possible (clauses (2c–d) and (3c–d)). In fact, clauses (2d) and (3d) mean that underspecified conditions can be resolved by the way certain *other* constituents update the discourse (e.g., $K_{\alpha'}$ can be resolved by the way K_{β} attaches to K_{α} , where $R(\alpha', \alpha)$ holds for some R). This is what happens in (28):

- (28) a. John took an engine from Avon to Dansville.
 b. He picked up the boxcar
 c. and took it to Broxburn.

The definite *the boxcar* in (28b) means that the DRS K_{b_p} , which appears in the SDRS that represents the presupposed content of (28b), contains the underspecified bridging relation $B = ?$ to some underspecified antecedent object $v = ?$. As in (14), *Narration* connects the asserted content of (28b), which we label b_a , to the content a of (28a). And as before, this imposes a spatio-temporal constraint, that the boxcar is in Dansville. Thanks to clause (3d), this additional information, that’s added to K_{b_a} via $Update_{\text{SDRT}}$, resolves $B = ?$ and $v = ?$ in K_{b_p} to $B = \textit{in}$ and $v = \textit{Dansville}$ respectively. Thus in this example, the resolution of content in the presupposition b_p is determined by the way the *asserted* content b_a attaches to the context a . We’ll see the formal details of this example in §5.3.

4.3 Some Remarks about Background

We proposed in §3 that the principal relation that binds presuppositions to the context is *Background*. This captures the intuition that they provide given information. And the DICE rules predict that *Background* is the principal relation, because generally, the constituent in the context or the presupposition itself is stative (cf. the rule **Background**).

Since *Background* plays such an important role in this theory, it is important to define the constraints it places on the contents of the propositions it connects. This is because these constraints affect the interpretation of presuppositions through $Update_{\text{SDRT}}$. Ultimately, *Background* ensures that presuppositions are constrained by more than informativeness and consistency.

We have already encountered the temporal constraint on *Background*, supplied by the **Axiom on Background** (Lascares and Asher 1993, Asher *et al.* 1996). There are three other constraints. First, all the so-called simple rhetorical relations (Asher, 1993)—such as *Background*, *Narration*, *Elaboration*, *Parallel* and others—are veridical: that is, the terms they relate are true:

- **Veridicality:** If R is simple, then $R(\alpha, \beta) \rightarrow (\vee\alpha \wedge \vee\beta)$

This guarantees that the two propositions *There is a King of France* and *he is bald* in the discourse (29), which features *Background*, are both true:

(29) There is a King of France. He is bald.

Second, both *Background* and *Narration* require a common topic. This explains why discourses like (30) sound odd without any special prior context.

- (30) a. ??Max walked in. Mary dyed her hair black.
 b. ??Max smoked a cigarette. Mary had black hair.

The final constraint is designed to account for the behaviour of pronouns in discourses that feature *Background*. Compare the text (31)—where (31b) connects to (31a) with *Background*—with (32), where the relation is *Narration*.

- (31) a. A burglar broke into Mary’s apartment.
 b. Mary was asleep.
 c. He stole the silver.
- (32) a. A burglar broke into Mary’s apartment.
 b. Mary discovered the break-in the next day.
 c. ??He stole the silver.

Asher *et al.* (1996) use (31) vs (32) to argue that if the discourse features a condition of the form *Background*(α, β), then objects introduced in α are available for pronominal reference in the proposition γ that comes after β . This isn’t the case for *Narration*(α, β). But *Background* is a coordinating relation, and therefore α itself isn’t available. So why are its objects possible antecedent’s to γ ’s pronouns?

This pronominal behaviour is captured in SDRT via the different kinds of topics that are constructed for Narrations and Backgrounds (Asher *et al.*, 1996). We concentrate here on the topics for Backgrounds. When *Background*(π_1, π_2) is inferred, this induces a structural modification to the SDRS. It creates what’s known as a *Foreground-Background Pair*, which contains (i) a constituent π_3 in which the information of both the background constituent π_2 and the foreground π_1 is ‘repeated’ (π_3 is the topic and its content amounts to $Update_{DRT}(K_{\pi_1}, K_{\pi_2})$), (ii) the SDRS which contains the constituents π_1, π_2 and the condition *Background*(π_1, π_2)

forms a constituent, which gets labeled (with, say, π_4), and (iii) a link between π_3 and π_4 is introduced, known as *FBP* (standing for *Foreground-Background Pair*): i.e., $FBP(\pi_3, \pi_4)$ is introduced into the SDRS. Thus whenever an SDRS contains the conditions given in (33), it also contains those given in (34):

$$(33) \quad \boxed{\begin{array}{l} \pi_1, \pi_2 \\ \hline \pi_1 : K_{\pi_1}, \pi_2 : K_{\pi_2} \\ \text{Background}(\pi_1, \pi_2) \end{array}}$$

$$(34) \quad \boxed{\begin{array}{l} \pi_3, \pi_4 \\ \hline \pi_3 : \text{Update}_{\text{DRT}}(K_{\pi_1}, K_{\pi_2}) \\ \text{FBP}(\pi_3, \pi_4) \\ \boxed{\begin{array}{l} \pi_1, \pi_2 \\ \hline \pi_1 : K_{\pi_1}, \pi_2 : K_{\pi_2} \\ \text{Background}(\pi_1, \pi_2) \end{array}} \end{array}}$$

Intuitively, the topic π_3 ensures that the ‘main story line’ content in K_{π_1} is still accessible. *FBP* is a subordinating relation. And so the available constituents to which future representations of clauses can attach are: π_3 , π_4 and π_2 . Since there are referents in the topic π_3 which denote objects that were introduced in π_1 (since π_3 ‘repeats’ the content of π_1), this makes these objects available for future anaphoric reference. This predicts, correctly, that the burglar is available for anaphoric reference in (31). In contrast, Narrations like that given in (32) don’t have this topic structure. Moreover, the constituent π_1 which represents the first sentence isn’t an available attachment site. So the pronoun *he* in (32c) can’t refer to the burglar, as required.

To handle examples like (30b), we impose the constraint on *FBP* formation that an *FBP* can be constructed only if the backgrounding constituent sets the stage for the foregrounded constituent. That is, there must be some thematic link between the two constituents that can either be determined already from the discourse context or from world knowledge. So in (30b), we can’t form an *FBP* because neither the context nor conventional knowledge have made clear any link between Mary’s hair being black and Max’s smoking a cigarette. Though we would like to have more to say about thematic continuity, at present we must leave the matter here.

We’ll see in §5.3 that the *FBP* topic structure plays an important role in constraining the interpretation of presuppositions. However, in many cases, it’s not essential to the analysis. So on occasion, we will omit it for the sake of simplicity.

5 The Pragmatics of Presuppositions

We now analyse some examples of how discourse is updated with presuppositions. We will start with simple sentences (i.e., sentences without scope bearing elements), and then move on to complex sentences and multi-sentence discourse.

5.1 Presuppositions in Simple Sentences


Consider (15) again.

(15) The King of France is bald.

We must build the SDRS conditions from the grammar, and then use SDRT update to build the final representation of the sentence. We will give an informal synopsis first. We saw in §4.1 that the grammar will produce two SDRSs for (15): one representing the presupposed content that there is a King of France, and the other representing the asserted content that he is bald. Since this is the first sentence in the discourse, the SDRS for the whole sentence must be updated with the asserted content first (by adding it to the empty set), since the presupposed content demands a (rhetorical) anaphoric link. Now the presupposed content must be attached to this asserted content with a rhetorical relation. DICE is used to compute this. The presupposed content is stative, and so the rule **Background** given in §4.2 will mean that the presupposed content is attached with *Background*. That is, the SDRS for (15) will amount to: That there is a King of France is *Background* to his being bald.

Background's topic structure is thus invoked in the analysis of (15), just as it is in any SDRS containing *Background*. This topic structure doesn't have a crucial effect here, largely because (15) is a simple sentence and doesn't feature the kind of discourse phenomena that are central to our concerns in this paper. However, we spell out the topic structure in this simple example, for the sake of completeness. In future analyses, in order to simplify matters, we may omit the topic structure, unless it plays an important role in producing the right semantic effects.

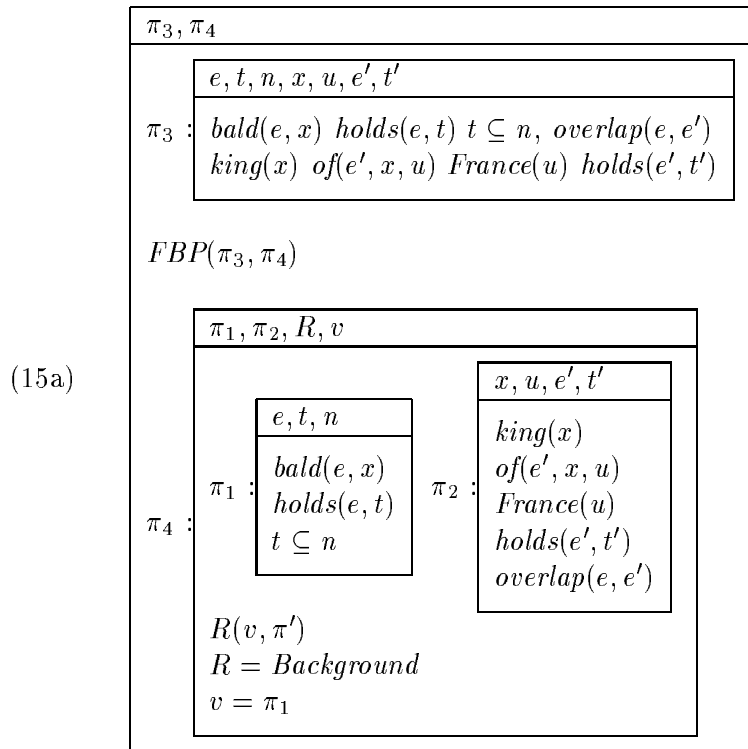
Let's now go into more detail. First, we construct the two SDRSs that the grammar produces for the presupposed content and the asserted content. The former includes a bridging relation *B* between the object denoted by the King of France and some (anaphoric) antecedent, that must be resolved. As argued in Asher and Lascarides (in press), the definite *the King of France* is a special case, where the grammar itself ensures that *France* and *of* serve to fill the underspecified antecedent object and bridging relation respectively.¹⁵ Therefore, the presupposed content of (15) that's produced by the grammar is the SDRS *p* below; its asserted content is *a*:



¹⁵This is a slightly simplistic analysis. For example, we have ignored the presuppositional content of proper names such as *France*, and we have also ignored any uniqueness conditions that might form part of the content of *The King of France*.

Now the task is to create an SDRS from these two constituents, using the definition of SDRT Update. (15) is the first clause in a discourse. So by clause 1 of the Update definition, the SDRS when the assertion a is processed is just a itself. We must then update this with p . According to clause 3 of SDRT Update, this involves using DICE to compute a rhetorical relation between π_2 and an available propositional discourse referent. In this case, there's only one: π_1 . Both π_1 and π_2 's main eventualities are states, and therefore **Background** applies. So by Defeasible Modus Ponens, $Background(\pi_1, \pi_2)$ is inferred. That is, that there is a King of France is background to his being bald. And according to clause 3 of the definition of SDRT Update, the underspecified conditions $R = ?$ and $v = ?$ need to be replaced with $R = Background$ and $v = \pi_1$ in the updated SDRS.

The condition $Background(\pi_1, \pi_2)$ imposes coherence constraints, described in §4.3. First, by **Axiom on Background**, $overlap(e, e')$ holds. So, by clause 3 of the SDRT Update definition, this condition must be added to K_{π_2} : that is, to the DRS that's labeled π_2 in p above. Second, the above topic structure is created, since the thematic link constraint is satisfied by the backgrounding constituent's giving a property of the subject of the assertion. So the final representation of (15) is (15a):



The SDRS (15a) entails that there is a King of France, because the axiom **Veridicality** applies to *Background* (and to *FBP*). Moreover, the King of France could be referred to with a pronoun in a subsequent sentence to (15), even though there being a King of France is background information. This is because π_3 and π_2 are available constituents.

This is a much more complex than the analysis given by van der Sandt (1992). Van der Sandt represents (15) as (15b), and this is identical to his analysis of (35):

x, u, e, t, e', t', n
<i>king-of-France</i> (x)
<i>bald</i> (e, x)
<i>holds</i> (e, t)
$t \subseteq n$

(35) A King of France is bald.

In fact, it's also (almost) identical to the topic π_3 in our representation. So we make the same predictions about which antecedents are available for future pronominal reference, in continuations of the discourse (15).¹⁶ Our analysis of (35) is (15b). This is different from the representation (15a) for (15). The difference arises because definites introduce anaphoric elements, including, for instance, the underspecified rhetorical relation R , which must be resolved by $Update_{SDRT}$. Indefinites aren't anaphoric in this way.

More importantly, however, our analysis of definites in discourse is more constrained than van der Sandt's. Van der Sandt imposes constraints such as informativeness, consistency, and accessibility. We strengthen these with pragmatics. To illustrate this, consider (36):

(36) ??A woman lives in that house. The man is bald.

According to our theory, the underspecified bridging relation B and the rhetorical relation R that's triggered by the content of the definite in (36) must be resolved through $SDRT$ Update. But B can't be resolved in this case, because reasoning about the rhetorical link fails to produce a plausible connection between the man and an antecedent object (the woman or the house) (see Asher and Lascarides, in press for details). In contrast, van der Sandt doesn't explain why (36) is odd: since there is no antecedent man to bind the presupposition to, one attempts to accommodate it, and since the result is consistent and informative, the accommodation is successful. So (36) is predicted on van der Sandt's account to be acceptable, when it is not.

5.2 Presuppositions in Conditional Sentences

Conditional sentences such as (1) and (2) give rise to choices about where to attach presupposed content, because the *Condition* relation, that's introduced by *if*, is a subordinating relation.

(1) If baldness is hereditary, then Jack's son is bald.

(2) If Jack has a son, then Jack's son is bald.

Like van der Sandt, the presupposed content in the consequent can attach in the global position (i.e., to the conditional as a whole), the intermediate position (i.e., to the antecedent)

¹⁶Of course Russell's analysis of (15) is different from ours and van der Sandt's because of his uniqueness condition, but we won't go into that here.

or the local position (i.e., to the asserted content of the consequent). Our glue logic must therefore yield inferences about which attachment site is chosen.

Recall Beaver’s (1995) idea that the interpretation of presuppositions is influenced by which reading is most plausible. One could imagine invoking here a general Gricean constraint of the form: make all of your attachments or presupposed and asserted information so as to make the speaker’s contribution as plausible as possible. We think, however, that the choice between attachment points that is governed by plausibility is specifically linked to conditional contexts and is a matter of discourse coherence. We claim that there is a linguistic convention about the strength of a Defeasible Consequence (or conditional) relation *vis a vis* a higher attachment with *Background* that is affected by plausibility. And it is just such a conventional comparison between the strength of discourse relations and attachments that **Maximize Discourse Coherence** can exploit. Indeed, the theory predicts differences between (10) and (11), because for (10) **Maximize Discourse Coherence** conflicts with the default **Prefer Global Attachment** and will therefore override it, whereas in (11) there’s no conflict:

- (10) If David wrote the article, then the knowledge that no good logician was involved will confound the editors.
- (11) If David wrote the article, then the knowledge that David is a computer program running on a PC will confound the editors.

SDRT alone has little to say about plausibility itself. In general, this is a matter of extra-linguistic knowledge, and even the task of comparing the plausibility of one scenario over another requires reasoning with arbitrary domain knowledge. Defining a logic of plausibility would therefore take us too far from our main concerns here. So, we’ll simply assume that we have some partial, qualitative or modal notion of comparative plausibility in DICE: $PL(\alpha) \succ PL(\beta)$ means α is more plausible than β .¹⁷

Let’s now look at conditional contexts in more detail. Given an attachment of $\partial(C)$ to B in the context $A > B$, where $>$ is some standard conditional operator, we would expect $A > (B \wedge \partial(C))$ to be equivalent to $A > B$ and $A > \partial(C)$. This equivalence holds in many conditional logics (for instance, Stalnaker’s (1968) conditional logic, Lewis’s (1973) counterfactual logics, or the normality conditional of Asher and Morreau 1991 which underlies DICE).

There are two consequences of this. The first is that the attachment of the presupposition $\partial(C)$ to the consequent B of a conditional whose antecedent is A , is equivalent to attaching $A > \partial(C)$ to the constituent in which the conditional $A > B$ occurs. So local attachments in this case are equivalent to conditional presuppositions. And if we assume that $>$ is the operator version of the relation *Def-Consequence*, then such local attachments in conditional contexts

¹⁷There are many ways one could develop this notion. For instance, one could tentatively identify the plausibility of a conditional with the conditional plausibility of the consequent upon the antecedent. If so, we’ll also want to assume that propositions can have a relative or conditional degree of plausibility $PL(\beta|\alpha)$, just as they can have conditional probabilities (in fact plausibility could be interpreted as subjective probability, though we will remain agnostic on this issue). On the other hand, one might explore a modal approach to comparative plausibility; $PL(\alpha) \succ PL(\beta)$ iff there are some α worlds that are closer to the expected worlds (given the content of the discourse so far) than any β worlds are. But again to pursue this discussion would lead us too far afield.

amount to attaching the presupposed information with *Def-Consequence*. The second is that such an attachment is plausible only to the extent that the conditional $A > \partial(C)$ is plausible. Of course we might attach presuppositions generated in the antecedent of a conditional to that antecedent with *Def-Consequence* as well; our observations about plausibility would hold of such attachments as well.¹⁸

It doesn't appear to be controversial or computationally onerous to assume that the information transfer function μ can tell us whether a conditional occurs in a constituent or not or whether one constituent is a subconstituent of another. We'll write $(\gamma \Rightarrow \alpha)(\pi)$, where $\gamma \Rightarrow \alpha$ is a gloss for a DICE predicate on the propositional variable π , to mean that $K_\gamma \Rightarrow K_\alpha$ is a DRS condition in K_π ; to capture subconstituency, we'll write $Sub(\alpha, \beta)$. Finally we'll say that one update of τ with β is strictly preferred to another when of the pair it is τ, β -maximal, which was defined in §4.2.

With this information in place, we now specify formally a constraint on discourse updates that will ensure that local or conditional attachments of presuppositions will be chosen only in cases when the plausibility of the resulting conditional is higher than the plausibility of the presupposition on its own. Further, the constraint tells us that normally when the conditional plausibility is higher, the conditional attachment is preferred. So Beaver's observations amount for us to the following: the presupposition $\partial(C)$ should attach locally (i.e., to B) normally if and (truth functionally) only if $A > \partial(C)$ is more plausible than $\partial(C)$ on its own (since global attachment entails $\partial(C)$).

¹⁸Some examples suggest that conditional presuppositions are not generated solely from the consequents of conditionals. Here is one due to Beaver (p.c.):

(37) If David is a shepherd and has a good rapport with his dog, he'll get a job on the farm.

This for at least some speakers generates the conditional presupposition:

(38) If David is a shepherd then he has a dog.

As it stands, our account could handle this example, provided that the plausibility of David's having a dog given that he was a shepherd is higher than his having a dog alone. Note, however, that it is the conditional that licenses this presupposition. The conditional presupposition vanishes when there is no conditional in context:

(39) David is a shepherd and he has a good rapport with his dog.

Further, the presupposition changes depending on the modal context. For example in (40) we only get the counterfactual presupposition in (41):

(40) If David were a shepherd and had a good rapport with his dog, he would get a job on the farm.

(41) If David were a shepherd then he would have a dog.

If the English indicative conditional is modelled by $>$, then the following theorem of the glue logic can be exploited:

$$((A \wedge B) > C) \leftrightarrow ((A \wedge B) > (B \wedge C))$$

Any presupposition that apparently surfaces from the antecedent of a conditional can also, in virtue of this equivalence, be interpreted as generated from the consequent. Thus, apparent examples of conditional presuppositions generated by the antecedents of conditionals can be interpreted as conditional presuppositions generated by the consequents, and we have already accounted for the latter. The story is the same for counterfactuals as in (40); the counterfactual conditional has a different semantics from $>$, but it still supports the equivalence above. In general we would like to thank David Beaver for the examples that led us to think about conditional presuppositions generated from the antecedents of conditionals.

- **Conditional Presuppositions:**

Suppose $Update(\tau, \alpha, \partial(\beta))$ links α and $\partial(\beta)$ with *Def-Consequence* and $Update(\tau, \alpha_1, \partial(\beta))$ links α_1 and $\partial(\beta)$ with *Background* and $(\alpha \Rightarrow \gamma)(\alpha_2)$ where $\alpha_2 = \alpha_1$ or $Sub(\alpha_2, \alpha_1)$. Then:

- $Update(\tau, \alpha, \partial(\beta))$ is strictly preferred to $Update(\tau, \alpha_1, \partial(\beta))$
 (normally if) and only if
 $PL((\alpha > \partial(\beta))(x))$ is greater than $PL(\mu(K_{\partial(\beta)})(x))$

This statement on preference of updates will play an important role in determining whether the presupposition should be locally or globally bound in conditional sentences such as (1), (2), (10) and (11).


Note that the above axiom on its own doesn't guarantee that a unique attachment site is chosen, even for conditional sentences. Heim (1992, 1993) and Van der Sandt (1992) note that presuppositions have a conventional property: they project out from embeddings to have the widest scope possible. In fact, in certain cases, this tendency to project actually seems to win over what's plausible. Consider sentence (11), for example. According to commonsense, it's more plausible that sentence (42) is true than (43):

- (42) If David wrote the article and he's a computer program, then the knowledge that David is a computer program will confound the editors.
- (43) David is a computer program, and if he wrote the article, then the knowledge that David is a computer program running on a PC will confound the editors.



(42) is the reading one gets of (11) if the presupposition attaches in the intermediate position. (43) is the result of global attachment, and corresponds to the intuitive interpretation of (11) (since it entails David is a computer program). So global attachment seems to be favored here, even though the result is the less plausible scenario. These examples show that, in contrast to Beaver (1996), a preference for attaching presuppositions that's based solely on plausibility cannot be right.

It's important to stress that (42) being dispreferred isn't a counterexample to the **Conditional Presuppositions** constraint. That axiom together with **Maximize Discourse Coherence** only rules out a conditional attachment, unless it's most plausible. But the presupposition in (42) isn't attached with a conditional or *Def-Consequence*, and so the constraint doesn't apply. There seems to be some tension between deriving a plausible reading, and the conventional property of presuppositions that they project to the widest scope. The constraint **Conditional Presuppositions** and the axioms of **Maximize Discourse Coherence** and **Prefer Global Attachment** reflect this tension. **Prefer Global Attachment** will determine that global attachment occurs in (1) and (11). But **Conditional Presuppositions** with **Maximize Discourse Coherence** will override this default to produce local attachment in (10).

With these constraints on conditional attachments in place, let's analyse (1) and (2). Let (1)'s antecedent that baldness is hereditary be represented by a (so K_{τ_1} supplies the information that baldness is hereditary):



The consequent (that Jack’s son is bald), contains a presupposition, and so it’s represented by two constituents b_a and b_p given below (as before, we have glossed over the presupposition triggered by the proper name *Jack*, by making the simplifying assumption that *Jack* is represented by the constant j):



The SDRS after updating with the first clause is simply a itself. Moreover, because of the cue word *if*, the grammar produces a *Condition* relation between the speech act discourse referent π_1 in a and another discourse referent, which must be equal to or include within its scope the asserted information b_a of the consequent. We will label the result—which includes the *Condition* relation, a and b_a —with the speech act discourse referent π_4 .

Now SDRS Update must be used to attach the rest of the discourse information, namely b_p , which is labeled π_3 . *Condition* is a subordinating relation, and so there are three available attachment sites for π_3 : either (a) π_3 binds ‘globally’, to the speech act discourse referent π_4 ; or (b) π_3 binds in the ‘intermediate’ position, to π_1 ; or (c) π_3 attaches in the ‘local’ position, to π_2 . One must check which of these choices produces a coherent result, and then choose among these coherent alternatives, using the axioms we specified earlier.


Attaching π_3 to π_1 isn’t possible: **Background** applies, but $Background(\pi_1, \pi_3)$ cannot be inferred, because a decent topic cannot be constructed; there is no thematic link between background and foreground, just as there is no thematic link in the incoherent (44).

(44) Baldness is hereditary. Jack has a son.

However, π_3 can attach to π_4 or π_2 coherently, both with *Background* (since **Background** applies in each case). We claimed earlier that the local attachment corresponds to the reading: If baldness is hereditary, then Jack has a son and he is bald. We can now explain why. The *Condition* relation holds between π_1 and the topic that’s introduced by the *Background*

relation between π_2 and π_3 , and this topic contains a ‘repeat’ of the content given in b_a and b_p . Thus part of the meaning of this SDRS is that if baldness is hereditary, then Jack has a son.

This dependency of truth, between the hereditary properties of baldness and Jack having a son, is relatively implausible compared with the reading where no such dependency is assumed. The latter reading is given by the global binding. Therefore, the axiom **Conditional Presuppositions** doesn’t apply. However, the default **Prefer Global Attachment** applies and so by **Defeasible Modus Ponens**, global binding is chosen in this case. The final analysis is given in (1’) below (in simplified form, since we have not included the *FBP* structure that’s introduced by *Background*):



Now consider the analysis of (2). As before, the grammar produces a constituent a which contains the information in the antecedent clause, which we label π_1 again. And it induces a *Condition* relation between this and (at least) the asserted information in the consequent. So just as with (1), there are three available attachment points for binding π_3 in the presupposed information b_p : π_4 , π_1 , and π_2 .

In (1), π_3 could not attach to π_1 with a rhetorical relation. But now the situation is different, because the content of π_1 is different. As we suggested in §3, if the content of the new information is a defeasible consequence of the content of the proposition to which it is to be attached, then normally, the rhetorical relation *Def-Consequence* is inferred. This is encapsulated in **Defeasible Consequence**:

- **Defeasible Consequence:**
 $((\tau, \alpha, \beta) \wedge ((\mu(K_\tau)(\tau) \wedge \mu(K_\alpha)(\alpha)) > \mu(K_\beta)(\beta))) > \text{Def-Consequence}(\alpha, \beta)$

The constraints on *Def-Consequence* are compatible with those on *Background*, and so it’s possible for both **Def-Consequence** and **Background** to fire, yielding the result that both *Def-Consequence*(α, β) and *Background*(α, β) hold. Note also that $\phi > \phi$ is valid for any proposition ϕ . So **Defeasible Consequence** will apply whenever one attempts to attach a presupposition to a constituent of the same content. Of course, our constraint on conditional presuppositions will hold here too. But in the cases of *Def-Consequence* where the axiom **Defeasible Consequence** is exploited, it only makes sense to suppose that the plausibility of the conditional (especially in this case since $\phi > \phi$ is a theorem) is always higher than the plausibility of ϕ alone (unless ϕ has the same status of certainty as a logical truth). So our constraint on **Conditional Presuppositions** and **Maximize Discourse Coherence** force us in the analysis of (2) to bind π_3 to π_1 (i.e., to bind the presupposition to the antecedent of

the conditional) with *Def-Consequence*. We therefore infer a binding to the (intermediate) attachment site with *Background* and *Def-Consequence*, even though the presupposition could coherently bind at the global level with *Background*. Thus **Maximize Discourse Coherence** and **Conditional Presuppositions** capture in a more general way van der Sandt’s principle that binding is preferred over accommodation. And for (2), these axioms predict, correctly, that the presupposition doesn’t project out from the conditional.

In general, whether a presupposition projects from an embedding or not depends on several things. First, it depends on pragmatic and semantic content of the prior discourse and the presupposition, because this is used to reason about which available attachment sites the presupposition can coherently bind to, and which rhetorical relation to use. Second, the projection depends on the relative plausibility of the various choices of attachment (that are coherent). Third, it depends on the relative strength of the rhetorical connections provided by various choices of attachment. Finally, it depends on the default which favors attaching presuppositions high up in the SDRS structure. Some general results concerning projection follow from this. Some of these are given below (for single sentences in the null context):

1. Projection does *not* occur in a sentence of the form *If A then B*, when the presupposition triggered by *B* is a default consequence of *A*.
2. Projection does occur in a sentence of the form *If A then B*, when the presupposition triggered by *B* is logically independent of that given by *A* and plausibilities reflect logical independence in the way that probabilities do (note that it can always attach coherently with *Background* to the global site, because *B* will contain sufficient content that’s similar to its presuppositions so that the topic required by *Background* can be constructed).
3. In a sentence of the form *A and not B*, the presuppositions triggered by *B* will project from the embedding so long as they can coherently attach to *A* with a rhetorical relation.

To say more in general is tricky, because projection is determined by a complex interaction between semantic and pragmatic knowledge resources. Furthermore, we view these solutions to the Projection Problem as a side issue; we’re more concerned with addressing the problem of interpreting presuppositions in multi-sentence discourse.

We now consider the analyses of (10) and (11). Van der Sandt’s analysis was inadequate, because it failed to distinguish their presupposed contents. In our theory, the constraint **Conditional Presuppositions** accounts for this variation. Consider (10) first. The presupposition trigger *knowledge that ϕ* generates a propositional presupposition ϕ that must be attached to the context. The conditional that if David wrote the article then no good logician was involved seems more plausible (given that in the context we know something about David) than the simple assertion that no good logician was involved.¹⁹ So local binding (to the asserted information in the consequent) is inferred for (10) via **Conditional Presuppositions** and **Maximize Discourse Coherence** in the now familiar way. In contrast, the conditional that if David wrote the article, then he’s a computer program is no more plausible than David’s

¹⁹Note that there is a logical dependence between the statements *David wrote the article* and *no good logician wrote the article*. And so the general rules we’ve stated about projection, which arise from $Update_{\text{SDRT}}$ and the plausibility axioms, predicts that the presupposition doesn’t project from the embedding.

being a computer program *tout court*. So the truth functional “only if” part of **Conditional Presuppositions** forces one to eschew the conditional attachment and to bind the presupposition in (11) non-locally (i.e., as *Background* to the antecedent, or to the whole conditional statement). **Prefer Global Attachment** leads us to infer global attachment, because of the default preference for it. Note that global attachment occurs even though the intermediate attachment describes a more plausible scenario, as we explained earlier. Our theory predicts this, because plausibility is a deciding factor only in quite specific circumstances (namely, when local attachment in the consequent of a conditional or when *conditional* attachment to the antecedent of a conditional is the most plausible). Thus the axioms which constrain the update of discourse with presuppositions predict the difference between (10) and (11).

It’s interesting to compare this approach to conditional presuppositions to that offered by satisfaction theorists (e.g., Beaver, 1997). Satisfaction accounts always generate conditional presuppositions in contexts such as (10). But such conditional presuppositions don’t always occur (cf. (11)). So they have tried to provide accounts for how the unconditional presuppositions arise from the conditional ones. Geurts (1996) argues convincingly that these attempts are flawed. This approach is different: the conditional presuppositions occur only when they meet certain plausibility constraints.

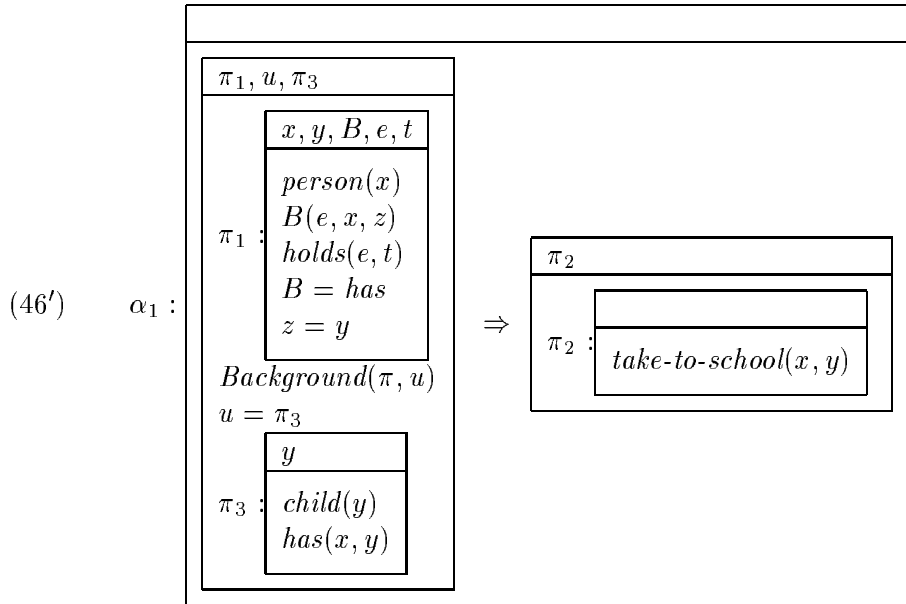
In §2, we observed examples where intermediate accommodation is predicted in standard theories of presupposition (e.g., van der Sandt, 1992), even though intuitively it should be blocked:

- (7) ?? I don’t know whether the Pope has measles. But every Catholic realizes the Pope has measles.

We now examine these examples in more detail. First, note that intermediate accommodation isn’t always blocked:

- (45) Nobody regrets leaving school.
 (46) Everybody takes their children to school.
 (47) Most people in the neighborhood take their children to school.

These cases apparently all involve information in the presupposition that serves to resolve the underspecified elements like the *B* relation of §4.1 (example 24) generated by the treatment of the generalized quantifiers. Further, it is just these resolutions that will allow us to infer an appropriate rhetorical link between the presupposition generated in the nuclear scope of the quantification and the restrictor. This predicts that intermediate attachment is allowed in (45-47). The presupposed material in these cases describe a property of the ‘quantified’ discourse referent. For instance in (46), the presupposed material is that *x* has children, where *x* is the discourse referent introduced by *everybody*. So we can set *B* to the “has” relation, and so relate *x* to *y* where *y* is a discourse referent representing a child of *x*, and at the same time establish the relevant thematic link between presupposed information arising from the nuclear scope and asserted information in the restrictor so as to establish a coherent *Background* relation (and *FBP* which we omit below for simplicity). So the representation of (46) is (46’):



In the cases where intermediate accommodation is blocked, the presupposed information is not of the right type to be identified with the contextually specified element in the restrictor of the quantifier. More specifically, it's not a property of the discourse referent that's being quantified over. So it doesn't give rise to an appropriate instantiation for the bridging relation B , and thus doesn't provide the appropriate thematic continuity for *Background*. So some and perhaps all of the underspecified conditions in the restrictor fail to receive a resolution. Consider for instance (7). The presupposed material that the Pope has measles is not a property of every Catholic. More precisely, it does not involve the variable introduced by *every catholic*. So the presupposed information cannot help resolve the bridging relation B or its other term. And we cannot get a thematic link so as to validate a *Background* relation between the presupposition and the restrictor; so the presupposition fails to bind to the intermediate position. As we explained in §2, the presupposition should not bind globally because of Moore's paradox. Assuming that the constraints on rhetorical relations capture this constraint on assertability, global attachment will be blocked in SDRT. In fact, local attachment will also be blocked in this way; note that (48), which is the reading given by local attachment, is also unassertable:

- (48) I don't know whether the Pope has measles. But for every person, if he's a Catholic, then the Pope has measles and he realizes the Pope has measles.

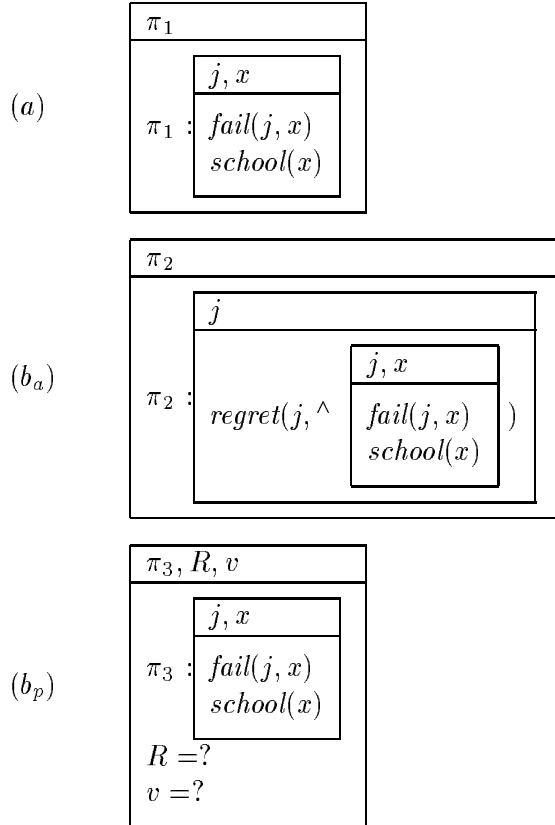
So SDRT predicts that (7) is odd.

5.3 Multi-sentence Discourse

In this section, we examine presuppositions in multi-sentence discourse. Consider (49):

- (49) a. John failed at school years ago.
 b. He now regrets that.

The pronoun *that* in (49b) introduces a propositional anaphoric discourse referent. One must resolve this to an available discourse referent of the same type. This gives only one choice: the proposition expressed by (49a). So, the pronoun *that* is identified with the content of (49a), regardless of how the constituents attach together with rhetorical relations. Therefore, the representation of (49a) is a , and the representation (49b) is b_a for the asserted content and b_p for the presupposed content (in slightly simplified form):



We now attempt to update a with b_p . **Defeasible Consequence** applies to π_1 and π_3 . And the temporal and topical constraints on *Background* are met, since π_1 and π_3 specify the same content. So, $Update_{SDRT}$ yields $Background(\pi_1, \pi_3)$ and $Def-Consequence(\pi_1, \pi_3)$ via DICE, and the condition $R = ?$ and $v = ?$ in b_p are replaced with $R = Background \wedge Def-Consequence$ and $v = \pi_1$.

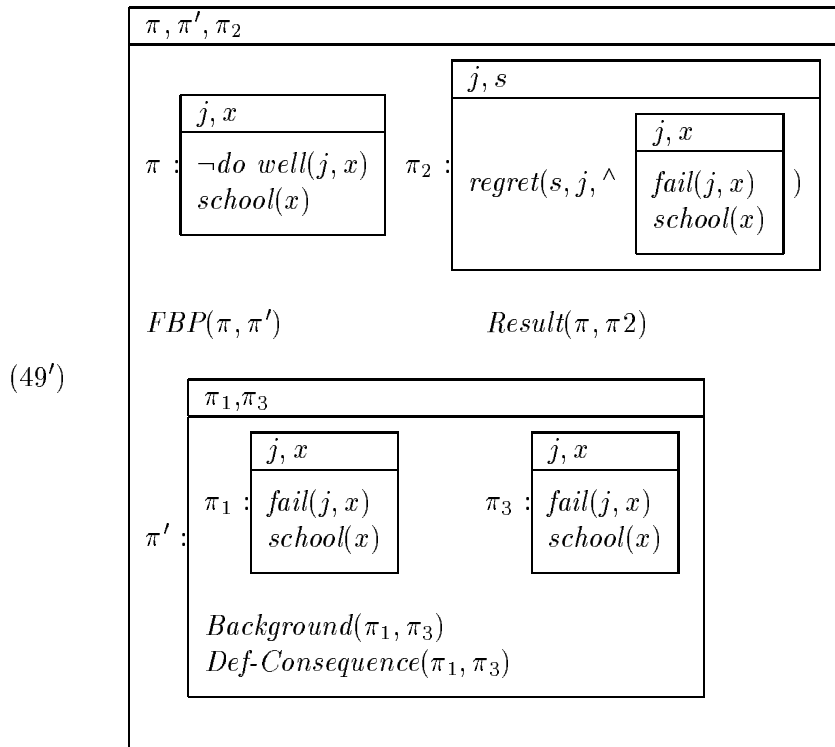
Furthermore, a topic π is built on top of this SDRS π' which contains a , b_p and the *Background* relation. The content in the topic π is a repeat of the content α and b_a , which is just a . And $FBP(\pi, \pi')$ holds. This topic structure is important in this case, because it means that the asserted content b_a can attach to the content of a , which is ‘repeated’ in π —for note that a itself is blocked from attachment by the *Background* relation with the presupposition. Intuitively, we want the asserted content of (49b) to be connected to the content (49a), because there is a causal relation between the action of failing in school and regretting it. The topic π lets us do this. In fact, since the presupposed content in this case is exactly that of (49a), attaching π_2 to the presupposition π_3 would have been adequate. But this would not be the case in the slightly modified example, given in (50):

- (50) a. John failed at school years ago.
 b. He now regrets that awful mistake.

Suppose *Background* didn't produce a topic that 'repeated' the content of (50a) and the presupposition (that there is an awful mistake). Then the regret would have to attach to the awful mistake with a rhetorical relation. So the regret would be caused by the awful mistake, rather than the action of failing in school *and* the fact that this was a mistake. By repeating all the content in the topic, new elements can attach to both the foreground and background information as required.

In (49), therefore, in order to capture this intuition that the regret is caused by the action described in (49a), we attach π_2 to the topic π with *Result*. Therefore, the final representation of (49) is (49'):

- (49) a. John failed in school.
 b. He now regrets that.



(49) is analogous to van der Sandt's *binding* of presupposed content. It's the limiting case, in which the presupposed information is a defeasible consequence of the constituent it attaches to, and it attaches with *Background* and *Def-Consequence*. In contrast, (51) is an example where the presupposition would traditionally be accommodated:

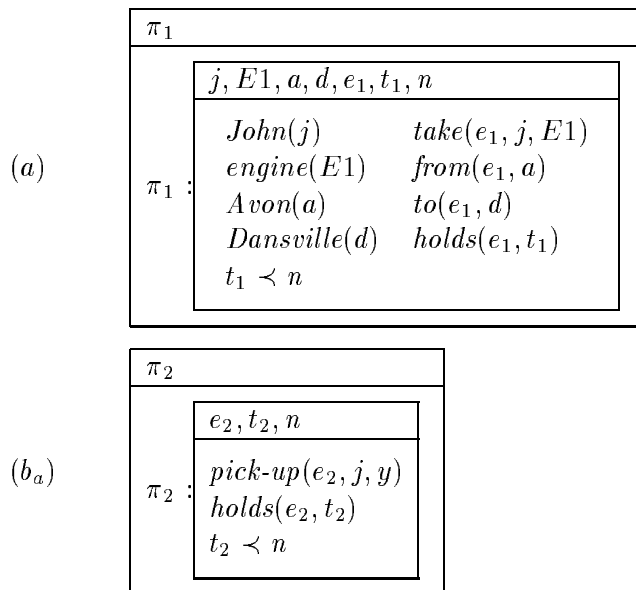
- (51) a. The US bombing of Baghdad during the Gulf War was much more indiscriminate and brutal than the American public was led to believe.
 b. But military commanders will not acknowledge that their campaign resulted in massive civilian casualties.

(51b)—which we label π_b —contains the presupposition trigger *acknowledge*, which introduces the presupposition $\partial(\pi_b)$ that their campaign resulted in massive civilian casualties. The cue phrase *but* indicates that the asserted content (that the military commanders will not acknowledge $\partial(\pi_b)$) must attach with *Contrast* to the discourse context (51a)—which we label π_a . This *Contrast* imposes coherence constraints: the two propositions π_a and π_b must have *contrasting themes* (Asher, 1993). In particular, the truth of one should lead to an expectation of the negation of the other. We must check this is the case. We must also compute a rhetorical attachment between $\partial(\pi_b)$ and an available attachment site. Suppose we attempt to attach $\partial(\pi_b)$ to π_a (i.e., attach ‘globally’). Then given that world knowledge supports a defeasible consequence relation between π_a and $\partial(\pi_b)$, they are connected with *Background* and *Defeasible Consequence* via Defeasible Modus Ponens on the rules **Background** and **Defeasible Consequence**. If $\partial(\pi_b)$ attaches this way, then the coherence constraints on *Contrast* between π_b and the (modified) discourse context—which contains π_a , $\partial(\pi_b)$ and a *Background* between them—are verified: there is a contrasting theme between $\partial(\pi_b)$ (which is in the discourse context) and *not acknowledge that $\partial(\pi_b)$* . The alternative to this is to attach $\partial(\pi_b)$ ‘locally’ instead of ‘globally’: i.e., to π_b itself. But this won’t produce as strong a contrasting theme. Therefore, the default preference for global attachment wins out by **Maximize Discourse Coherence**, and ensures that $\partial(\pi_b)$ attaches to π_a rather than π_b , because this produces the better discourse. In sum, **Background as Defeasible Consequence** ensures that *Background*($\pi_a, \partial(\pi_b)$) holds; this produces a *FBP* topic π which contains the semantic content of π_a and $\partial(\pi_b)$; and then π_b attaches to π with *Contrast*.

Consider now an example which involves bridging as well as presupposition satisfaction:

- (28) a. John took an engine from Avon to Dansville.
 b. He picked up the boxcar
 c. and took it to Broxburn.

(28a) is represented as a , while the asserted and presupposed components of (28b) are b_a and b_p :



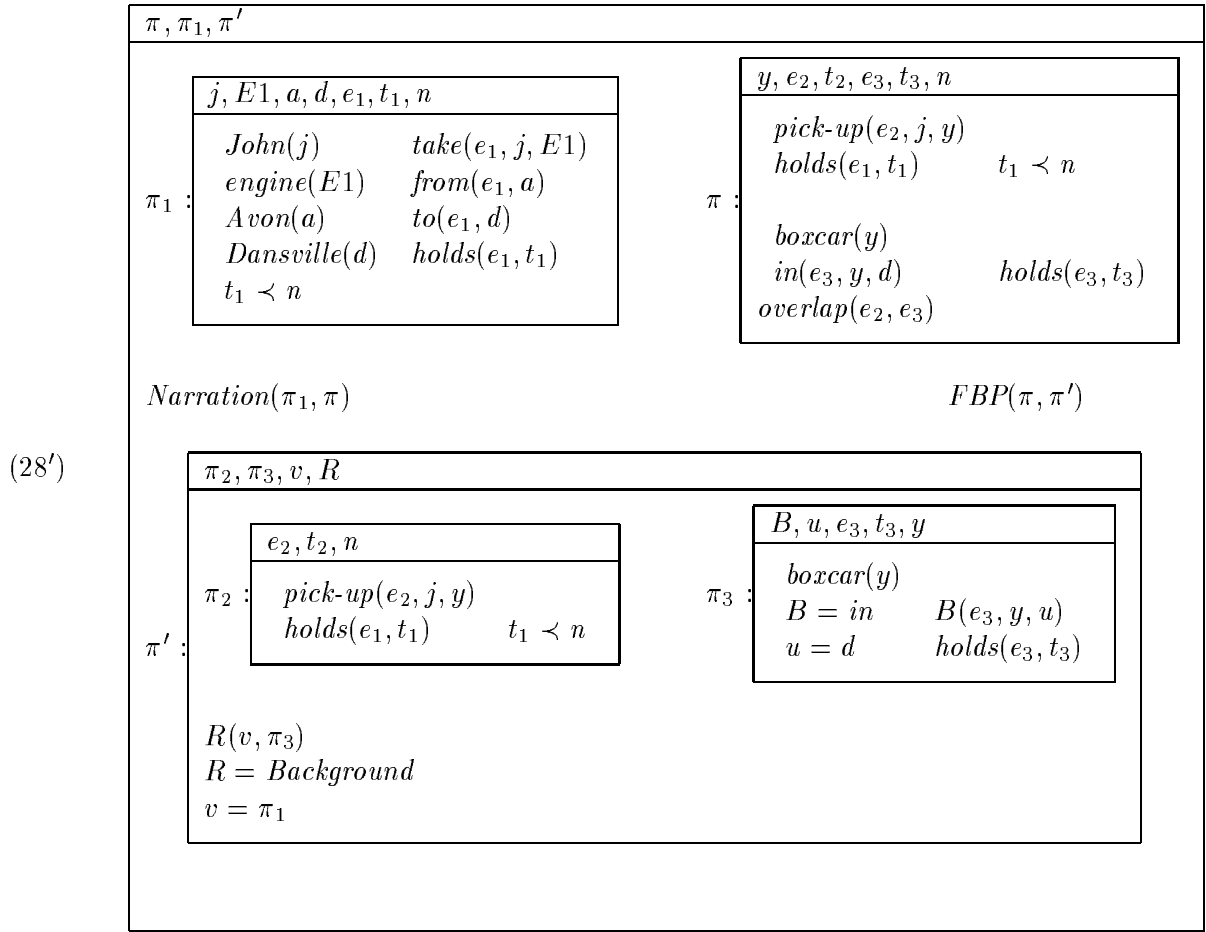
π_3, v, R									
(b_p)	$\pi_3 :$								
	<table style="width: 100%; border-collapse: collapse;"> <tr> <td colspan="2" style="padding: 2px;">B, u, e_3, t_3, y</td> </tr> <tr> <td style="padding: 2px;">$B = ?$</td> <td style="padding: 2px;">$B(e_3, y, u)$</td> </tr> <tr> <td style="padding: 2px;">$u = ?$</td> <td style="padding: 2px;">$holds(e_3, t_3)$</td> </tr> <tr> <td colspan="2" style="padding: 2px;">$boxcar(y)$</td> </tr> </table>	B, u, e_3, t_3, y		$B = ?$	$B(e_3, y, u)$	$u = ?$	$holds(e_3, t_3)$	$boxcar(y)$	
	B, u, e_3, t_3, y								
	$B = ?$	$B(e_3, y, u)$							
$u = ?$	$holds(e_3, t_3)$								
$boxcar(y)$									
$R(v, \pi_3)$									
$R = ?$									
$v = ?$									

According to $Update_{SDRT}$, we should check whether there is sufficient information in π_1 , π_2 and π_3 that we can attach them together with rhetorical relations, which lead to the underspecified elements in π_3 being resolved. Suppose we were to try and attach π_3 to π_1 (or π_2) first: π_3 cannot attach to π_1 (or π_2) with any relation—in particular *Narration* or *Background*—because there is insufficient information in π_3 to trigger the relevant DICE axioms.

So let's ignore π_3 for now, and try to attach π_2 to π_1 , to see if this produces possible resolutions of underspecified elements in π_3 .²⁰ And indeed, it does. π_1 and π_2 both describe events, and so Defeasible Modus Ponens on the axiom **Narration** yields $Narration(\pi_1, \pi_2)$. We must then accommodate content arising from *Narration*'s coherence constraints. First, e_{π_1} occurs before e_{π_2} : that is, the taking of the engine from Avon to Dansville occurs before y (which is the boxcar according to K_{π_3}) is picked up. Second, by modus ponens on **Spatial Consequence on Narration**, the actor *John* is in the same place at the end of taking the engine to Dansville as when he starts to pick up y . By the prepositional phrase in (28a), this is Dansville. And by the lexical semantic content of the phrase *picking up*, this event starts in the same location as it finishes, and the object y that's picked up is also at this location. Therefore, y is in Dansville.

Thus, the coherence constraints on *Narration* provide a particular way of resolving some of the underspecified conditions in the presupposed information—viz. B resolves to *in*, and u to *d* or Dansville. So the definition of $Update_{SDRT}$ ensures that B and u are resolved this way. So the content of b_p now contains sufficient information for us to compute a rhetorical relation between it and the discourse context: since B has resolved to *in*, e_3 is a state. Moreover, attaching π_2 to π_1 with *Narration* has made only π_2 available for attachment. Therefore Defeasible Modus Ponens on **Background** yields $Background(\pi_2, \pi_3)$. The temporal constraint on *Background* means that e_{π_2} and e_{π_3} overlap (i.e., picking up the boxcar and the boxcar being in Dansville temporally overlap). This is as required. Moreover, a topic is computed by creating a new constituent π with all the content of b_a and b_p repeated in it, and this is related via *FBP* to the SDRS π' , which contains b_a , b_p and the condition $Background(\pi_2, \pi_3)$. Because this new topic π gets built over the SDRS π' , the original constituent π_1 that was attached with *Narration* moves up to be attached to π (for details, see Asher 1993 and Asher *et al.* 1996). The final result is given in (28'):

²⁰Note that clauses (2c-d) and (3c-d) in $Update_{SDRT}$ allow content φ , that's inferred by attaching π_2 to π_1 , to affect the resolution of underspecification in π_3 to occur, because π_3 will attach to π_2 or π_1 .



Note that we predicted in this analysis that the presupposition is locally accommodated to π_2 : π_2 must be attached first because π_3 isn't informative enough, and since the rhetorical relation is *Narration*, π_1 is blocked from future attachments. This predicts correctly that the engine in (28a) cannot be an antecedent to a pronoun in a continuation of (28): e.g., *?John drove it quickly*. However, *John drove the engine quickly* is a much more acceptable continuation. And our theory reflects this: *the engine* requires a bridging relation, and computing the bridging relation is possible, by adding to π_2 (and hence to the topic π) the content that the original engine in (28a) (to which π_2 is linked) is used to pick up the boxcar.

Let us now return to an example which was problematic for van der Sandt's account:

- (4) Either John didn't solve the problem or else Mary realizes that the problem has been solved.

The presupposition that the problem's been solved cannot bind anywhere coherently, because unresolved bridging relations would remain, which are triggered by the definite *the problem*. So sentence (4) is predicted to be odd in this null context.

However, in the appropriate discourse context, (4) is coherent:

- (6) a. A: The Problem Solving Group is given a problem each day, and the group leader Mary has to assign it to someone in her group. John is the best problem solver. But when he solves a problem, he always boasts about it. This annoys Mary, and so if she thinks that the day's problem is an unsolved one, she gives it to him, to test him. Otherwise, she gives it to another member of her group.
- b. B: John's being very quiet just now. Did she give him today's problem?
- c. A: Well, I'm not sure she did. Either John didn't solve the problem or else Mary realized that the problem's been solved.

The overall structure of this example is that (6b) is a response to (6a), perhaps some sort of *Commentary*, while (6c) is a response to the question in (6b) (we'll ignore the internal structure of (6b) here). The disjunction in (6c) offers two, contrasting *Explanations* to the response *I'm not sure*. The explanations explain why John is behaving quietly (i.e., why (6b) is a fact).

How do the presuppositions fit into this structure? *A* sets up a relationship between Mary and John in (6a); Mary gives problems that she believes are unsolved to John, and problems that she believes are solved to the others. (6b) refers to a problem (*today's problem*), to which the presuppositions triggered by the definite *the problem* in (6c) can bind. Note that (6b) is available to these for binding, because (6c) is a response to (6b). So the bridging relation and antecedent are resolved to identity and today's problem respectively.

The second sentence in (6c) triggers two further presuppositions. *John didn't solve the problem* presupposes he got the problem, and *Mary realized that the problem's been solved* presupposes that the problem's been solved. Note that given the content of (6a), this second disjunct also implies, at least defeasibly, that John didn't get the problem.

Let's look at the binding possibilities for these presuppositions. Binding *John got the problem* locally with *Background* allows the coherence constraints on *Contrast* to be satisfied: the first disjunct now implies that John got the problem, and the second implies that he didn't. It also ensures that the second sentence in (6c) does provide two contrasting *Explanations* to the answer in the first sentence of (6c): either (i) John did get the problem because Mary believed it was unsolved, but he couldn't solve it, or (ii) he didn't get the problem because Mary believed the problem already had a solution (note that *realize* entails *believes*). In effect both (i) and (ii) allow us to compute answers to *B's* question as to whether Mary gave John the problem: (i) implies that he got the problem (and presumably that she gave it to him); (ii) implies that she didn't. They both also offer contrasting explanations for what happened yesterday, and their disjunction explains why *A* isn't sure about how to answer *B's* question.

Now consider what happens if we attempt to bind the presupposition that John got the problem globally. In that case, the explanation for the first sentence to (6c) becomes: John got the problem, and either he didn't solve it, or Mary realizes that the problem's been solved. This utterance implicates that John got the problem, but from someone other than Mary, because the second disjunct implies that Mary didn't give the problem to John. This utterance would form a coherent explanation for John's quiet behavior only if it's interpreted as follows: John got the problem from someone other than Mary, and either he didn't solve it, or he did, but because Mary realized the problem had been solved and so John didn't get the problem

from her, he didn't boast that he has solved it. Thus, there are more assumptions required to support this explanation than for the local binding case. Intuitively, the explanation for the local binding case is stronger than this one. But then our principle of **Maximize Discourse Coherence** predicts that the presupposition of *John didn't solve the problem* is bound locally.

Now consider the presupposition that the problem's been solved. According to the discourse structure, *the problem's been solved* could bind locally to *Mary realizes that the problem's been solved*, to the first sentence *I'm not sure she did* in (6c), or to the question to which (6c) is a response, namely the utterance (6b). In this case again, the constraints on discourse relations and discourse update in SDRT determine a preferred attachment site, and once again **Maximize Discourse Coherence** plays a crucial role. The *Contrast* between options (i) and (ii) is strengthened if the presupposition of *realize* binds locally to *Mary realizes the problem's been solved* with *Background*, rather than being bound globally. To see this, consider the interpretation of the discourse under local binding, *vs.* that given by global binding. The 'local' interpretation is: Either (i) John got the problem because Mary believes it was unsolved, but John didn't solve it, or (ii) John didn't get the problem because the problem has been solved and Mary realized this. The 'global' interpretation is: The problem's been solved, and either (i) John got the problem because Mary thought it hadn't been solved, and he failed to solve it, or (ii) John didn't get the problem because Mary realized that the problem's been solved.

In the local case, *Contrast* relates two propositions where one contains the content that John didn't solve the problem and the other contains the content that the problem has been solved. But the 'global' case doesn't relate two propositions with this content. Rather, the *Contrast* in the global case is supported only by the contrast that is there for the local case too, between John getting the problem and not getting it. So, the *Contrast* relation is better in the local case than in the global one. Therefore, **Maximize Discourse Coherence** yields the prediction that the presupposition is bound locally in this discourse, rather than globally.

In contrast to van der Sandt (1992) and to Gazdar (1980), we have predicted local attachment in (6), because of SDRT's constraints on discourse relations and discourse update. The context in (6) doesn't implicate that the problem has been solved; nor does it implicate that it hasn't been. And therefore one cannot determine that the presupposition should be bound locally in these theories in virtue of the conversational implicatures supported by the context, at least in any straightforward interpretation of Gazdar's views. Given what *A* said in (6a), *A* and *B* would know the answer to the question as to whether John got today's problem, if they knew if Mary believed it was solved. This context implicates that *A* and *B* don't know if Mary believes the problems been solved. However, as we explained earlier, the context does *not* implicate whether or not *A* and *B* themselves believe that the problem's been solved. And so these implications on their own don't determine whether the presupposition should be bound globally or locally. Hence, the coherence constraints on *Contrast*, and the preference for interpreting discourse in a way that strengthens rhetorical connections, both play a crucial role in interpreting the presuppositions in this example. This demonstrates that constraints imposed by discourse structure are in general stronger than informativeness and consistency, and indeed go beyond constraints provided by conversational implicature alone.

5.4 Presuppositions, Cognitive States and Global Accommodation

We have sketched already in §5.1 how the constraints on *Background* rules out many cases of infelicitous accommodation: the presupposition to be accommodated and the global attachment site don't have the requisite thematic continuity. However, what appear as global accommodations often seem felicitous. Consider the following from Vallduvi (1990):

- (52) a. A: I bought the President a tray for his Delft china set.
b. B: That wasn't a good idea. The president [hates]_F the Delft china set.
(from Vallduvi (1990))

We claim that in (52) the presupposition introduced by the focussed phrase *hates* is bound by means of *Background* to a belief of the interpreter about the speaker. That is, the presupposition binds the presuppositions to what one reconstructs about the speaker's cognitive states. To reconstruct the *Elaboration* or *Explanation* given by *B*'s response to *A*, we need to introduce a wide range of assumptions about what *B* thought was behind *A*'s utterance. Ordinary theories of focus linked with a theory of presupposition (e.g. Rooth 1992) claim that some element of the alternative set for the focussed expression is presupposed. But the standard view of the presuppositions of focus would in this case simply accommodate this element, when the presupposition is in fact *bound* to what *B* reconstructs about *A*'s mental state from what *A* said.²¹

The possibility that presuppositions can bind via *Background* or even *Def-Consequence* to an interpreter's beliefs to be sure puts some slop into our discourse based account of presupposition. Nevertheless, in contrast to the standard account, we predict that such so called global accommodations obey more constraints than consistency and informativity. Our approach would in principle distinguish between (53) and (54) (Asher and Lascarides, in press):

(53) The present King of France does not exist.

(54) ?The man does not exist.

In contrast to (15), *Background* cannot hold between the asserted and presupposed content in (53), because the two states of the king existing and not existing cannot overlap. We claim that *Background* holds in (53) between the presupposed information and the cognitive state of the hearer, who is familiar with the concept, the King of France. But as (54) shows, this strategy is not available in all cases. With (54) there simply isn't enough content associated with the description to provide the necessary thematic continuity and to bind the presupposition to any particular concept via *Background*.

But even in cases where this strategy isn't available, it appears that sometimes presuppositions can be introduced felicitously in information that is 'discourse structure initial' (i.e., is simply added and not attached to an empty discourse context). Consider the first sentence from Austen's *Sense and Sensibility*:

²¹For some details as to how this story should go, see Asher (1995).

The family of Dashwood had been long settled in Sussex. Their estate was large, and their residence was at Norland Park, in the center of their property, where for many generations they had lived in so respectable a manner as to engage the general good opinion of their surrounding acquaintance.

Austen knows that there isn't any Dashwood that the reader knows about, but she fills in the scene during the course of the novel. There is presupposition failure but the reader doesn't really care. He expects the author to tell him more later. Presupposition in such cases is not anaphoric but cataphoric upon what will come next. The licensing conditions for such cataphors is a matter for future research.

6 Conclusion

We have offered a new picture of presuppositions: we have modified van der Sandt's view that they are anaphoric by loosening the notion of binding, thereby foregoing the need for accommodation. We have demonstrated that there are three main advantages to processing presuppositions through binding them to the context with a rhetorical relation. First, accommodation by addition as in van der Sandt (1992) fails to account for the interaction that many have observed between the interpretation of presuppositions, domain knowledge and the content of the discourse context (e.g., Beaver 1996, 1997, Lascarides and Oberlander 1993). In contrast, binding material together with rhetorical relations involves reasoning about domain knowledge and semantic content, and so in essence, we have provided pragmatic constraints on presupposition satisfaction that is lacking in current dynamic accounts of presuppositions. Second, through taking rhetorical constraints into account, we model how the semantic content itself of a presupposition is dependent on pragmatic information (e.g., (10) vs. (11)), and we model how the level of accommodation (i.e., global vs. intermediate vs. local) is also pragmatically determined. The result is a more fine-grained analysis of presuppositions than has been offered so far. Finally, and arguably most importantly, we have reduced the problem of presupposition satisfaction to the problem of resolving scope ambiguities—a phenomenon that arises in many other areas of linguistic study.

Acknowledgements

This work is sponsored by a grant funded by ESRC UK (grant number R000236052), and by the ESRC-funded Human Communication Research Centre, University of Edinburgh. We would like to thank David Beaver, Hans Kamp, Rob van der Sandt, and an anonymous reviewer for helpful comments on previous drafts.

References

Asher, N. [1993] *Reference To Abstract Objects in Discourse*, Kluwer Academic Publishers.

- Asher, N. [1995]** From Discourse Micro-structure to Macro-structure and back again: The Interpretation of Focus, in H. Kamp and B. Partee (eds.) *Conferences on Semantics and Context*, Prague and Bad Teinach, SFB Report, University of Stuttgart, 1995.
- Asher, N. [1996]** Mathematical Treatments of Discourse Contexts, in P. Dekker and M. Stokhof (eds.) *Proceedings of the Tenth Amsterdam Colloquium on Formal Semantics*, pp. 21–40, ILLC publications, University of Amsterdam.
- Asher, N. [forthcoming]** The Logical Foundations of Discourse Interpretation, in J. M. Larrazabal (ed.) *Logic Colloquium 1996*, Springer Verlag.
- Asher, N., M. Aurnague, M. Bras, L. Vieu [1996]** De l’espace-temps dans l’analyse du discours”, *Semiotique : Numero Special "Theories sémantiques et modalisation*, **9**.
- Asher, N. and T. Fernando [1997]** Effective Labeling for Disambiguation, *Proceedings of the Second International Workshop in Computational Linguistics*, Tilburg, the Netherlands.
- Asher, N., D. Hardt and J. Busquets [1997]** Discourse Parallelism, Scope and Ellipsis, in *Proceedings of Semantics and Linguistic Theory (SALT VII)*, Stanford Ca, March 1997.
- Asher, N. and A. Lascarides [1994]** Intentions and Information in Discourse, in *Proceedings of the 32nd Annual Meeting of the Association of Computational Linguistics*, Las Cruces, USA, June 1994, pp34–41.
- Asher, N. and A. Lascarides [1995]** Lexical Disambiguation in a Discourse Context, *Journal of Semantics*, **12** number 1, pp69–108, Oxford University Press.
- Asher, N. and A. Lascarides [1996a]** Questions in Dialogue, Research report number HCRC/RP–81, *Linguistics and Philosophy*, in press.
- Asher, N. and A. Lascarides [in press]** Bridging, to appear in *Journal of Semantics*.
- Asher, N. and Morreau, M. [1991]** Common Sense Entailment: A Modal Theory of Non-monotonic Reasoning, in *Proceedings to the 12th International Joint Conference on Artificial Intelligence*, Sydney Australia, August 1991.
- Beaver, D. [1996]**, Presupposition, in J. van Benthem A. ter Meulen (eds.), *Handbook of Logic and Linguistics*.
- Beaver, D. [1997]**, Presupposition in DRT, *SALT VII*, Stanford, CA.
- Chierchia, G. [1995]** *Dynamics of Meaning : anaphora, presupposition, and the theory of grammar*, University of Chicago Press.
- Clark, H. [1977]** Bridging, in P. N. Johnson-Laird, P. C. Wason (eds.) *Thinking: Readings in Cognitive Science*, pp. 411–420, Cambridge University Press.
- Fernando, T. [1994]** What is a DRS?, in R. Cooper and J. Groenendijk (eds.) *Integrating Semantic Theories II, Dyana Deliverable R2.1.B*, 1994, also appeared in the Proceedings of the First International Workshop on Computational Semantics, Tilburg 1994.
- von Fintel, K. [1994]** *Restrictions on Quantifier Domains*, PhD Thesis, University of Massachusetts, Amherst Mass.

- Gazdar, G. [1979]** *Pragmatics, Implicature, Presupposition, and Logical Form*, Academic Press, London.
- Geurts, B. [1995]** *Presupposing*, PhD. dissertation, Universität Stuttgart.
- Geurts, B. [1996]** Local Satisfaction Guaranteed, *Linguistics and Philosophy*, **19**, pp.259-293.
- Groenendijk, J. and M. Stokhof [1991]** Dynamic Predicate Logic, *Linguistics and Philosophy*, **14**, pp39–100.
- Grosz, B. and C. Sidner [1986]** Attention, Intentions and the Structure of Discourse. *Computational Linguistics*, **12**, 175–204.
- Heim, I. [1982]** *The Semantics of Definite and Indefinite Noun Phrases*, PhD thesis, University of Massachusetts, Amherst.
- Heim, I. [1983]** On the Projection Problem of Presuppositions, *Proceedings of the West Coast Conference on Formal Linguistics*, **2**, pp114–126.
- Heim, I. [1992]** Presupposition Projection and the Semantics of Attitude Verbs, *Journal of Semantics*, **9** (3).
- Heinamaki, O. [1972]**, *Semantics of English Temporal Connectives*, Ph.D. Dissertation, University of Texas at Austin.
- Hobbs, J. [1985]** On the Coherence and Structure of Discourse. Report No. CSLI-85-37, Center for the Study of Language and Information, October, 1985.
- Hobbs, J. R., M. E. Stickel, D. E. Appelt, P. Martin [1993]** Interpretation as Abduction, *Artificial Intelligence*, **63**, pp69–142.
- Kamp, H. and U. Reyle [1993]** *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*, Kluwer Academic Publishers.
- Karttunen, L. [1974]** Presupposition and Linguistic Context, *Theoretical Linguistics* **1**, pp. 181–194.
- Lascarides, A. and N. Asher [1993]** Temporal Interpretation, Discourse Relations and Commonsense Entailment, in *Linguistics and Philosophy*, **16**, pp.437–493.
- Lascarides, A., A. Copestake, and E. J. Briscoe, [1996]** Ambiguity and Coherence, *Journal of Semantics*, **13.1**, pp41–65, Oxford University Press.
- Lascarides, A. and J. Oberlander [1993]** Temporal Connectives in a Discourse Context, in *Proceedings of the European Chapter of the Association for Computational Linguistics (EAACL93)*, pp260–268, Utrecht, The Netherlands.
- Lewis, D. [1973]** *Counterfactuals*, Blackwells, Oxford.
- Lewis, D. [1979]** Scorekeeping in a Language Game, *Journal of Philosophical Logic*, **8**, pp. 339–59.

- Mercer, R. E. [1985]** *A Default Logic Approach to the Derivation of Natural Language Presuppositions*, PhD dissertation, University of British Columbia.
- Mercer, R. E. [1988]** Solving some Persistent Presupposition Problems, in *Proceedings of COLING 1988*, pp420–425.
- Moore, G.E. [1912]** *Ethics*, London.
- Muskens, R. [1995]** Order independence and Underspecification, in Dyana-2 Deliverable R2.2c: *Ellipsis, Underspecification and Events in Dynamic Semantics*, available from ILLC, University of Amsterdam.
- Polanyi, L. [1985]** A Theory of Discourse Structure and Discourse Coherence. In *Papers from the General Session at the Twenty-First Regional Meeting of the Chicago Linguistics Society*, pp. 25–27.
- Pollard, C. and I. A. Sag [1994]** *Head Driven Phrase Structure Grammar*, University of Chicago Press.
- Reyle, U. [1993]** Dealing with Ambiguities by Underspecification: Construction, Representation and Deduction, *Journal of Semantics*, **10.1**, pp123–179, Oxford University Press.
- Russell, B. [1905]** On Denoting, *Mind*, **14**, pp. 479–493.
- van der Sandt, R. [1992]** Presupposition Projection as Anaphora Resolution, *Journal of Semantics*, **19**(4).
- Stalnaker, R. [1974]** Pragmatic Presuppositions, in M. K. Munitz and P. K. Unger (eds.) *Semantics and Philosophy*, New York University Press, New York, pp197–213.
- Stalnaker, R. [1978]** Assertion, in Cole, Peter (ed.) *Syntax and Semantics volume 9: Pragmatics*, Academic Press.
- Thompson, S. and W. Mann [1987]** Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, **1**, 79–105.
- Vallduví, E. [1990]** *The Informational Component*, PhD. thesis, University of Pennsylvania, Philadelphia.