Questions in Dialogue*

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

May 3, 2001

Abstract

In this paper we explore how compositional semantics, discourse structure, and the cognitive states of participants all contribute to pragmatic constraints on answers to questions in dialogue. We synthesise formal semantic theories on questions and answers with techniques for discourse interpretation familiar from computational linguistics, and show how this provides richer constraints on responses in dialogue than either component can achieve alone.

1 Introduction

Understanding dialogue involves a complex interaction between the semantics of the utterances, the discourse context of which they are part and the cognitive states of the participants. In an effort to begin to spell out these connections precisely, we focus here on the role of questions and answers in dialogue—more specifically, *how*-questions. Like many researchers (e.g., Hobbs 1985, Polanyi 1985), we represent a discourse context as a structured object consisting of questions and propositions related by discourse relations. *How*-questions yield an intuitive connection between a particular sort of discourse structure and the intentions of the participants, since answers to a particular type of *how*-question are closely related to plans.

Grosz and Sidner (1990), Litman and Allen (1990), Pollack (1986, 1990), and others exploit domain knowledge about how to achieve actions to build the structure of the dialogue. This is an essential component of the task. But these theories haven't linked their accounts to the role of compositional semantics in building a discourse structure. We think this poses several problems: it forces them to rely on domain knowledge for information about participants' intentions in certain cases where this isn't plausible; and it fails to fully analyse the pragmatic/semantic interaction—in particular, how discourse relations can affect information about *both* the participant's intentional states *and* the semantic content of the discourse—thus missing important generalizations about dialogue.

The exact nature of discourse structure is determined in large part by what information one wants to extract from it during dialogue processing. This in turn determines which knowledge

[1]We are grateful to Jonothan Ginzburg, Rich Thomason and two anonymous reviewers for providing very helpful comments on previous drafts of this paper. Any shortcomings and mistakes that remain are entirely our responsibility.

representations are relevant for building it. We will argue that if discourse structure is to constrain the possible antecedents to anaphora, enable us to specify the semantic content of a sentence in the light of the way it connects to the previous discourse, and inform us about how the participants' cognitive states change as the dialogue proceeds, then domain knowledge about plans is an essential source of information for building discourse structure, but it's insufficient on its own. One needs insights from formal semantics too.

Segmented Discourse Representation Theory (SDRT) is a theory of discourse semantics designed to explore systematically the interface between compositional semantics, pragmatics and discourse structure. To date it has been used to model the effects of discourse structure on pronominal and temporal anaphora in monologue (e.g., Asher 1993, Lascarides and Asher 1993). Asher and Lascarides (1994) examine how discourse structure, the participant's intentions and semantic content interact in monologue. But SDRT still lacks an account of how discourse structure interacts with plans. However, we will argue that it is well-placed to model the interaction between discourse structure and cognitive states, and therefore provides a good forum in which formal semantics, discourse structure and cognitive states can be folded together. We will combine the existing techniques in SDRT for encoding the interaction between formal semantics and discourse structure, with techniques for plan recognition. The result will be an extension of SDRT which achieves an analysis of dialogue with the advantages of insights from both the theoretical linguistics literature on the compositional semantics of questions and answers, and the computational linguistics literature on plan recognition. We will demonstrate how this works by analysing in detail *how*-questions in cooperative dialogue. To motivate our analysis, we first consider some simple examples.

2 Motivating Examples

In this section, we consider some interrogatives and responses in simple dialogues, and discuss what an adequate analysis of these dialogues must achieve. For the sake of clarification, we make a terminological distinction between questions as syntactic objects and questions as semantic and illocutionary objects. The former are *interrogatives*, and the latter are *questions*. Similarly, the utterance in a dialogue that addresses a question is a *response*, and utterances that *semantically* speaking are answers to the question are called *answers*.

Very simple examples of interrogatives and responses in dialogue reveal the complex interaction between discourse structure, participants' mental states and semantic content.

(1) a. A: How can I get to the treasure?
b. B: It's at the secret valley.
c. A: But I don't know how to get there.

An adequate analysis of (1) must provide answers to at least the following questions: How does A infer from (1b) that she should go to the secret valley? On what basis is (1b′) an odd response to (1a)? And why is the use of *but* coherent in (1a,b,c), but not in (1a,b,c′)?

(1) b′. B: ?Mary's hair is black.
c′. A: ?But I know how to get there.

On an intuitive level, answers to these questions involve reasoning about compositional semantics, discourse structure and cognitive states. For instance, the compositional semantics and discourse structure must represent the fact that (1a) is a question, (1b) is a response to it, and *it* in (1b) refers to the treasure. Moreover, a combination of the compositional semantics of (1b) (from which one learns that the treasure is at the secret valley) and A's cognitive state (more specifically, his knowledge about action and causation), should be sufficient to infer at this stage that A would achieve getting to the treasure by executing the action of going to the secret valley. This must contrast with the semantic contribution of (1b′), where intuitively A's cognitive state doesn't combine with the information in (1b′) to help A infer what he has to do to get to the treasure. On the other hand, the semantics of (1c) and the way it attaches to (1a,b) in the discourse structure should allow B to infer that A still can't execute a plan of getting to the treasure, even though he knows it's at the secret valley. Typically, B must then revise his model of A's cognitive state (i.e., revise his assumption that A can execute a plan to go to the secret valley).

We believe that an adequate analysis of dialogues like (1) lies outside the scope of both formal semantic and extant AI theories of dialogue. Plan recognition techniques alone don't provide a precise semantics of cue words like *but* and so they don't detect a difference in acceptability between (1c) and (1c′). One exception is Green and Carberry (1994), but their plan-based semantics of *but* requires a taxonomy of plans and plan operators that isn't motivated independently of dialogue analysis. Consequently, providing a general axiomatisation of how one reasons with the various types of plans required is problematic. On the other hand, formal semantic theories of interrogatives such as Groenendijk and Stokhof's (1984) provide some epistemic notions of answerhood; they use these to define what it means to give an answer to a question. But their accounts do not invoke the complex cognitive states or discourse structures that, as argued by Bromberger (1962), are necessary for analysing the semantic content of *how*- and *why*-interrogatives and the pragmatic constraints on answerhood.

Dialogue (2) below exemplifies a truism that is nevertheless difficult to capture with extant theories of dialogue: people often use the compositional semantics of an utterance to learn different ways of doing an action, when the first attempt to do it failed.

(2)
a. A: How do I install the assembly?
b. B: Typically, you put the assembly on before tightening the screws on the assembly.
c. A: But I tried that and it was too difficult.
d. B: So let's tighten the screws on the assembly first then.

The domain knowledge in (2b) is ultimately explicitly ruled out as being irrelevant to achieving the action. If one cannot use a detailed semantics of the utterances to construct the discourse structure, then it's unclear how to predict in precise ways exactly how the participants' domain plan for installing the assembly can change as the dialogue proceeds. B's suggestion in (2d) need not be based solely on an existing plan in his knowledge base. Rather, he has actively constructed it via the semantics of the dialogue. Discourse structure, though sometimes built up by exploiting domain knowledge, also often affects and even disregards domain knowledge. So one cannot use domain knowledge alone to build plans for executing speech acts or discourse structure; formal semantics is a crucial source of information too.

Dialogues (3) and (4) show that constraints on interpretation given by compositional semantics, in this case concerning anaphora, can refine those given by intentional structure.

(3)
a. A: How can I get to 6th Street?
b. B: You can get there by asking someone Downtown.
c. A: ?What's his name?
c′. A: Oh yeah? And could you tell me his name?
c″. A: Who exactly should I ask? A passerby?

(4)
a. A: How can I get to 6th Street?
b. B: There's someone Downtown that you could ask.
c. A: What's his name?

Intuitively, one cannot use the pronominal reference *his* in (3c) in (3a,b,c), whereas it's acceptable in (3c′) in (3a,b,c′) and in (4c). It is difficult to see how the plan recognition techniques for dialogue analysis could predict this. This is because the domain level plan alluded to in both (3) and (4) are extremely similar, if not exactly the same. In both cases, the relevant plan is to ask directions for 6th Street from someone who has a sufficient knowledge of the local area. It is conceivable that one could obtain a distinction in the domain level plans along the following lines: that the one underlying (3) involves going Downtown and choosing the person to ask *randomly*, whereas with (4) the person isn't chosen randomly. But it is difficult to motivate this, because it requires an axiomatisation over very fine grained distinctions among the plan operators during dialogue processing. It would put a huge load on this axiomatisation. Moreover, assuming that (3b) and (4b) have different underlying speech acts also lacks independent motivation, especially if B's cognitive state in each case is the same. For example, it's perfectly possible that in both (3b) and (4b), B has a *de dicto* belief that someone knows the way. For even when B utters (4b) rather than (3b), B need not have a particular person Downtown in mind, that you could ask. Therefore, recognising the speech acts doesn't distinguish (3) and (4) either. At any rate, reasoning about the participants' cognitive states is in general difficult and uncertain, and so it would be better to predict constraints on pronominal reference from monotonic procedures for building compositional semantic representations whenever it is possible to do so.

Here, the difference between (3b) and (4b) lies in the scope relations between the existential quantifier (*someone*) and the modalities (*know* and *can*). The difference between (3c) and (3c′) lies in the fact that they connect in different ways to the context (3a,b); there's some aspect of *contrast* in connecting (3c′) to (3b), which is lacking in (3c), and there is also a modal subordination in (3c′) that is lacking in (3c). This contrast relation and modal subordination results in different constraints on the available antecedents to anaphora (Asher, 1993).[1] We need to make use of a semantics in which these scope differences in (3b) vs (4b), and the rhetorical differences in (3a,b,c) vs (3a,b,c′) affect anaphoric possibilities.

More generally, the surface similarities between (3b) and (4b) and (3c) vs (3c′) show how small changes in the surface structure of a sentence should sometimes have profound effects on discourse structure, so that the changes in anaphoric possibilities that result from these small changes in the sentences are recorded. We aim to constrain the kinds of content that are coherent in follow-up utterances in dialogue (so asking a question about who to ask after hearing (3a,b) or (4a,b) is perfectly admissable, for example, whereas talking about the color of Mary's hair would sound odd). That is, we wish to constrain *what* can be said next. But we also aim to go further, and constrain *how* the thing to be said can be worded. These latter constraints on linguistic formation of content should account for when pronouns can be used and when they can't, for example. Using compositional semantics to build discourse structure helps achieve this. In essence, we aim to use compositional semantics to explain why asking who to ask is expressed adequately in (3c′) and (3c″), but not in (3c).

3 Previous Approaches to Dialogue

Many researchers in AI have demonstrated that plans are an important source of information about discourse structure (e.g., Grosz and Sidner 1990, Litman and Allen 1990, Green and Carberry 1994). For example, plans can be used to construct a discourse structure of an appropriate kind in the task of dialogue generation (e.g., Hovy, 1988). However, discourse *interpretation* involves intention recognition, and plans are not always the most useful clue for inferring the discourse structure of the dialogue and its semantic content. Moore and Pollack (1992), for example, demonstrated that there isn't a one to one mapping between dialogue and plan structures.

[1] Note that replacing *And* in (3c′) with the explicit contrastive item *But* doesn't change the overall content of this utterance, thereby providing evidence that (3c′) contrasts in some way with the context.

Here, we take this further and argue that when the task at hand is discourse interpretation as opposed to generation, a semantic-based theory of discourse structure like SDRT is needed to assess when exploiting plans is appropriate.

Grosz and Sidner (1990) and Litman and Allen (1990) emphasise the distinction between a domain level plan and a discourse level one. In other words, one can intend to do things in the domain—such as eat dinner—and one can intend to perform a speech act—such as persuade the interpreter that a particular proposition is true. The two structures aren't always isomorphic. For example, in dialogue (5), at the domain level, the relevant plan is to catch the train.

(5)
- a. A: I need to catch the 1:20 to Philadelphia.
- b. Where's it leaving from?
- c. B: Platform 7.
- d. A: Where do I get a ticket?
- e. B: From the booth at the far right end of the hall.

One can assume that by world knowledge, this plan has the structure represented below: catching a train involves (or, in Grosz and Sidner's terminology, *dominates*) two subactions; buying a ticket and going to the platform, where the former satisfaction-precedes the latter (assuming that you must have a ticket before you board the train):

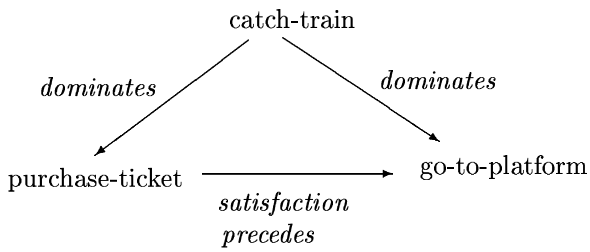

On the other hand, at the speech act level, A has requested information about where to buy a ticket *after* he requests information about the platform, and so the former speech act doesn't satisfaction-precede the latter. Hence the plan for performing speech acts in (5) isn't isomorphic to the domain level plan.

Grosz and Sidner (1990) and Lochbaum (1995) rightly argue that commonsense plans and discourse level ones are closely related, and are often isomorphic. They use techniques discussed by Perrault (1990) to compute the speech act that underlies the utterance currently being processed. Then, using a conversational default rule, they infer from it that the participants want to achieve a particular shared plan. This plan is partially specified, and filling it out involves complex inference about domain knowledge and what gets said in the discourse. They rely on commonsense plans from the plan library, which encapsulates knowledge about how to do actions, to reason about how this shared plan relates to previous ones that were inferred from previous utterances via the conversational default rule. The resulting intentional structure is then used to structure the dialogue into segments. They argue that the natural segmentation of a discourse is isomorphic to this intentional structure. In other words, a segment that's embedded in another in the linguistic structure means that the discourse segment purpose (DSP) or speech act of the former segment is dominated by that of the latter in the intentional structure. And when one segment follows another, the DSP of the first satisfaction-precedes that of the second.

In the particular case of dialogue (5) this appears to yield unintuitive results, when attaching (5d) to the preceding discourse. Applying their principles in the most straightforward way produces the following segmentation of the discourse segment (5a,b,c).

a — A: I need to catch the 1:20 to Philadelphia.
b — Where's it leaving from?
c — B: Platform 7.

This is because the conversational default rule yields from (5a) that A and B intend to achieve a shared plan that A catch the 1:20 train to Philadelphia; and from (5b) it yields that A and B intend to achieve a shared plan that B inform A which platform the train is leaving from; (5c) continues this discourse segment purpose (DSP); and in the above domain level plan, the dominance relation between catching a train and going to the platform allows one to conclude that the DSP for (5a) dominates that for (5b,c). So we get an embedding in the discourse segmentation as shown above, because of the isomorphism between linguistic structure and intentional structure.

Processing (5d) proceeds along similar lines, and the conversational default rule is used to infer that A and B mutually believe that A desires to achieve a shared plan that B informs A of where the ticket booth is. According to the above commonsense plan, locating the ticket booth is dominated by catching the train. So as before, one can use this aspect of the domain level plan to infer that (5d)'s DSP is dominated by that of (5a), and hence (5d) is embedded in the top level discourse segment in the above linguistic structure. However, using the above commonsense plan in this way to reason about how (5d)'s DSP is related to that of (5b,c) will produce a counterintuitive linguistic structure. Locating the ticket booth satisfaction-precedes going to the platform in the commonsense plan. If we continue to assume that the relation between DSPs are determined by the structural relations between the actions in the commonsense plan, then (5d)'s DSP will satisfaction-precede (5b)'s, and consequently, by the isomorphism of intentional structure and linguistic structure, the linguistic structure is one where (5d) is not on the right frontier; rather (5b-c) are:

(5a–d)
(5d) (5b–c)

This yields the wrong predictions about possible antecedents to anaphora for utterances that continue the dialogue (for Grosz and Sidner assume a right frontier constraint, in line with Polanyi (1985), Webber (1991) and others).

One might think that this counterintuitive result could be avoided, if instead of appealing to the structural relationship between the *actions* in the commonsense plan about catching trains, one appealed to the dominance and satisfaction-precedence relations between *knowing* how to catch a train, *knowing* where to buy a ticket, and *knowing* which platform the train is leaving from. In such a plan, knowing p dominates knowing q, if knowing p entails knowing q. On the other hand, to preserve a uniform interpretation of *satisfaction precedence* for epistemic level plans and domain level ones, it is necessary to assume that if knowing p satisfaction precedes knowing q, then there is no logical entailment relation between knowing p and knowing q; but it is necessary to know p before you can know q. Therefore, although knowing how to catch a train dominates knowing where to buy a ticket and knowing where the platform is, there are no satisfaction-precedence relations between the latter two pieces of knowledge. It isn't necessary to know where the platform is before you know where the ticket booth is; or vice versa. So if the relationship between (5a), (5b) and (5d) was determined by this epistemic level of the commensense plan about trains, rather than by the domain level actions themselves, then one wouldn't obtain a counterintuitive relationship between (5b) and (5c) in the discourse structure.

However, this way out of the problem won't work in general, because the order in which one knows things doesn't usually matter. Dominance relationships at the epistemic level are commonplace, because dominance is inherently linked to semantic entailment between the propositions that are the objects of the knowing relation. But satisfaction precedence relationships are rare at the epistemic level. It is difficult to think of cases where knowing one thing satisfaction precedes knowing something else, when there is no logical entailment relation between the two. There are some examples in mathematical proofs where this occurs, but this is quite a specialised context. Consequently, there isn't much structure to the epistemic level of a commonsense plan. And so appealing to this alone won't always provide sufficient information on which to determine discourse structure. In essence, satisfaction precedence would have to take a back seat in Grosz and Sidner's theory, thereby weakening their constraints on antecedents to anaphora.

Indeed, in some cases where plan structure can be used to explain a dialogue phenomenon, something else could be used instead. For example, one could explain the incoherence of (6) via reasoning about the commonsense plan.

(6) a. A: I need to catch the 1:20 to Philadelphia. Where is it leaving from?
b. B: Platform 7.
c. A: ?Where can I get information about trains to Philadelphia?

Regardless of whether one uses the underlying commonsense plan at the action level or epistemic level, the DSP underlying (6c) is associated with an item in the commonsense plan that dominates the DSP underlying (6a,b). Under the assumption that a latter utterance dominating a former one is an illegal discourse structure, the isomorphism between intentional structure and discourse structure predicts that (6) is odd. But although the incoherence of (6) can be explained by appealing to plan structure, it needn't be. It could also be explained by appealing to Gricean principles. A's interrogative in (6a) implies that he knows something about where to get information about trains to Philadelphia. He has asked B for this information, indicating that he believes he can get this information from B. In this context, A's interrogative (6c) sounds odd, because B can assume that A knows an answer to this question; namely, ask B. So A violates Gricean constraints on contributions in dialogue, by asking a question that B can presume A already knows the answer to. Note how A's utterance is much improved in (6a,b,c′):

(6) c′. A: Where can I get more information about trains to Philadelphia?

This explanation of (6) doesn't appeal to plans about how to catch trains at all, but instead appeals to Gricean-style pragmatic maxims.

Commonsense plans are an essential source of knowledge in dialogue processing. But we need a way of deciding when we can use the commonsense plan to reason about how DSPs are related to each other, and when we can't. This is currently lacking in Grosz and Sidner's (1990) and Lochbaum's (1995) analyses. We need to predict automatically when the intentional structure of the discourse is isomorphic to the commonsense plan, and when it isn't. By giving a formal theory of the integration between plans, constraints on discourse structure and compositional semantics, we can come closer to this goal.[2]

Intuitively, compositional semantics familiar from current formal linguistic theories could contribute to this decision about the relevance of the commonsense plan in structuring discourse. In particular, the *compositional semantics* of (5a) informs us that this sentence describes A's intention (via the use of *need*), and that of (5b) informs us that this is a question.[3] This semantic information, which can be built via *monotonic* procedures, could be used instead of the above domain plan structure, to infer that A asked the question because he doesn't have enough information to achieve the intention. In other words, the monotonic semantic information that is

[2] One could test the following hypothesis: the intentional and commonsense plan structures match as long as it doesn't violate constraints on discourse structure, such as the most recent utterance being on the right frontier.

[3] Also note that the dynamic semantics helps resolve *it* in (5b) to the train (cf. the account of anaphora resolution in Kamp and Reyle, 1993).

retrieved from what is made linguistically explicit is used to nonmonotonically infer, via the fact that the intention and the question are in juxtaposed utterances, that the intention prompted the question. Thus (5b) is seen as subordinate to (5a) in the discourse structure because we reasoned via semantics that the question was prompted by the intention, rather than reasoning that one subordinates the other because this is what happens in the domain structure.

However, the commonsense plan would be useful in constraining cooperative responses to (5b), given the context in which A asked it: Any coherent response should elaborate a plan that will enable A to achieve the intention that prompted the question in the first place. So a coherent answer to (5b) must elaborate a plan for catching the 1:20pm train to Philadelphia; thus one must check that (5c) elaborates the above plan for catching a train. A similar line of reasoning would underly the analysis of (5d) and its response (5e)—dynamic semantics, rather than plan structure, would be the central clue for reasoning about how (5d) relates to (5a) and (5c), but then the commonsense plan must be used to ensure (5e)'s an appropriate response, given A's purposes. We will spell out this analysis precisely in section 6.2.

Grosz and Sidner's analysis of dialogue would benefit from a precise formal representation of semantics, since semantic information could help one select exactly when the commonsense plans are relevant to structuring dialogue, via a model of inference that links cognitive states, discourse structure, *and* semantics. The same holds for all AI theories of dialogue that use plan structures to interpret discourse. There are four problems that remain because such inferences aren't modeled. First, one can't attack discourse segmentation. Since discourse segmentation is inferred via plan structure, we need inferences about which plan is most plausibly the one underlying the discourse, so that we can rank choices. This leads to the second problem. Since a lack of inference procedure precludes one from choosing among alternative intentional structures, one can't motivate the discourse structure constructed in dialogue analysis, except by theory bound intuitions. Thirdly, without a model of inference from discourse structure to semantic effects, one can't use the discourse structure to learn more about the semantic content of segments in the dialogue, such as the temporal order of events described in dialogues such as (7), or the inference that the House Bill 1711 is bad for big business in (8) (taken from Moore and Pollack 1992).

(7) A: John's late for work.
B: He must have missed the bus.
C: That was because he had to help with the baby.
B: Yeah, his wife hasn't been feeling well lately

(8) George Bush supports big business. He's sure to veto House Bill 1711.

4 Our General Strategy

We will use Segmented Discourse Representation Theory (SDRT) to combine compositional semantics, discourse structure and information about the participants' intentional states in order to over come the inferential problems encountered in the AI plan recognition work. We'll do this for simple dialogues involving the question/answer alternation, with the hope that this yields a general methodology for dealing with other patterns of dialogue.

SDRT has four main advantages. First, the way discourse structure affects and is affected by semantic content has already been studied extensively in this framework (e.g., Asher 1993, Lascarides and Asher 1993). Second, the basic semantic framework which underlies SDRT (Discourse Representation Theory or DRT, Kamp and Reyle 1993), specifies constraints on anaphora resolution that are sensitive to the relative scope of quantifiers such as *someone* and modalities such as *know* and *can*, and we argued earlier that this would be helpful for distinguishing the acceptability of (3) vs (4). Third, DRT has already been used to model cognitive states (e.g., Asher 1986, 1987, Kamp 1990, Singh and Asher, 1993), and so SDRT is well suited for folding discourse structure, semantic content and mental states together, so that coherence constraints on dialogue exploit information from all these levels. Finally, a distinctive feature of our analysis is that the rhetorical relations which are used to represent a dialogue impose constraints on multiple levels, not only on the semantic content of the constituents being connected together, but also on the discourse structure surrounding those constituents and on the cognitive states of the participants.

Our plan of attack is first to specify a formal semantics of interrogatives and responses in DRT. Second, we will extend SDRT so that it represents dialogue as well as monologue. We will then use the formal semantics of interrogatives and responses in DRT to help constrain answers to questions in dialogue in SDRT.[4] SDRT will impose constraints on the information flow between discourse structure, semantics and information about the intentional states of the participants in the dialogue. These will augment the existing coherence constraints already specified in the theory. For example, the new constraints on questions and answers and the existing constraints on *Contrast* (and hence the cue word *but*), will predict why (1a,b,c) is better than (1a,b,c′).

4.1 A Crash Course in SDRT

SDRT, which is an extension of DRT, is a theory of discourse structure and content. An NL discourse is represented by a segmented DRS (SDRS), which is a recursive structure of DRSs that represent clauses with rhetorical relations linking them. The rhetorical relations are modeled after those proposed by Hobbs (1985), Polanyi (1985), Thompson and Mann (1987) and others, and include *Explanation*, *Contrast* and *Continuation*, among others. To make technical definitions about SDRT below more tractable, we will adopt the labelling techniques for SDRSs from Asher (1996, forthcoming); that is, each SDRS will have a label and rhetorical relation symbols will take labels of SDRSs as constituents. We make the general notational convention that π labels the SDRS K_π. Using labels in the arguments to rhetorical relations will allow us to simplify the notion of SDRS update—that is, the process of extending an SDRS with new information—and it will also enable us to compare it with the simpler DRS update notion, as we'll see shortly.

As in DRT, SDRT uses the level of semantic representations, viz. that of SDRSs, to model the dynamic aspects of discourse interpretation. The way this is done is to define an update function on semantic representations of a given context and information to be integrated into that context. The update function returns, if successful, a new representation—or a new model of a discourse context within which further information can be integrated. Now, as shown in Asher (1993) one can give a completely model-theoretic interpretation of bottom-up DRT, and as in Muskens (1993) and Fernando (1994), one can prove that a particularly simple update function on DRSs defines an equivalent transition on information states construed as model embedding function pairs. This update function when applied to DRSs K and K' is simply defined as:

$$Update_{drt}(K, K') \quad = \quad (U_K \cup U_{K'}, C_K \cup C_{K'})$$

That is, one takes the DRS K which represents the discourse context, and the DRS K' which represents the new information in the current clause, and one represents the new discourse context as a DRS which is the union of the discourse referents in K and K', and the union of their conditions.

The update function for DRT becomes problematic as soon as we countenance conditions in DRSs that reflect unresolved anaphoric pronouns—of the form $x = ?$. To resolve these conditions, the update function must make an appeal to various sources of information, and some work is required to replicate Muskens's and Fernando's result (for details see Asher 1996, forthcoming).

[4]We encode these constraints in SDRT rather than the simpler DRT, because SDRT contains essential information about discourse coherence constraints on relations like *Contrast* (and hence cue words like *but*), which is missing in DRT.

Computing the update of a discourse is more complex in SDRT than it is in DRT; it's more complex even than DRT with ambiguous pronouns. SDRT exploits a number of information sources to define an appropriate *Update* relation. SDRSs are built up incrementally, and building an SDRS from NL discourse involves (i) choosing an attachment point in the discourse structure built so far at which to attach the new information, and (ii) computing which rhetorical relation should be used to attach the representation of the current clause to that attachment point. SDRT countenances a special logic for computing a rhetorical relation. We can think of this as a 'glue' logic or a logic of information packaging, by means of which we arrive at a more completely specified representation of the content of the discourse.

The glue logic differs from the logic of the "information content", which operates on the representations themselves and whose consequence relation is defined relative to the model-theoretic structures for those representations. Asher (1996, forthcoming) offers several versions of SDRT with several different versions of the logic of information content, following the approach to the logic underlying dynamic semantics developed in Fernando (1994). But these logics have a validity problem that is at least recursively enumerable (r.e.), for the expressive power of the SDRT language themselves is at least that of first order logic.

The glue logic, on the other hand, exploits a glue language, which is decidedly less expressive than that of SDRT itself. In earlier papers (e.g., Lascarides and Asher, 1993), it sufficed to make the glue language a quantifier free fragment of a first order language augmented by a weak conditional operator $>$, which formalizes generic or defeasible rules of interpretation ($P > Q$ means "if P then normally Q").[5]

In SDRT, the glue logic known as DICE (Discourse in Commonsense Entailment) exploits a variety of information sources—formal semantic and lexical information principally, as well as information about the domain, the cognitive state of the speaker, and about the conventions of language use. But these information sources are typically couched in much more expressive languages that the language of DICE. To exploit the information in SDRSs in DICE, we therefore make use of an information transfer function μ from the SDRS representation itself and from other richer languages (such as the language of the participant's cognitive states) into the weaker language of DICE, and this allows us to infer rhetorical relations using compositional semantic content and the interpreter's other knowledge resources as clues. For example, the transfer function μ, formalized in Asher (forthcoming), associates the SDRS labels that feature as arguments to rhetorical relations with information in the SDRS that the label refers to; roughly, it takes conditions inside the SDRS and turns them into predicates of the labels.

The DICE rules about which rhetorical relation holds exploit these predicates on labels. The rules themselves are typically default rules expressed using the connective $>$ and they represent the normal rhetorical role that a certain piece of information plays in discourse. These default rules enable the interpreter to come to nonmonotonic conclusions about which rhetorical relation to use to bind the constituents in an SDRS together. DICE also contains axioms concerning the semantic effects of rhetorical relations. These rhetorical relations and their semantic effects, that are inferred via DICE, are then used to define the appropriate update relation in SDRT.

As we've mentioned, the rules in DICE are default conditionals of the form $P > Q$. For example, `Background` below states: If β is to be attached to α with a rhetorical relation, where α is part of the discourse structure τ already (i.e., $\langle\tau,\alpha,\beta\rangle$ holds), and β describes a state (i.e., $state(e_\beta)$ holds), then normally, the rhetorical relation is *Background*.

- `Background:` $(\langle\tau,\alpha,\beta\rangle \wedge state(e_\beta)) > Background(\alpha,\beta)$
- `Axiom on Background:` $Background(\alpha,\beta) \rightarrow overlap(e_\alpha,e_\beta)$

This captures the rhetorical role that states normally play in discourse, as argued for in Lascarides

[5] For details, see Lascarides and Asher (1993).

and Asher (1993) and elsewhere: they normally provide background information. And typically, information like $state(e_\beta)$ will be part of the interpreter's stock of information expressed in the glue language on the basis of the transfer function μ, which enables DICE to access semantic information in (the more expressive) SDRSs.

All rhetorical relations are accompanied by coherence constraints. For example, the `Axiom on Background` states that constituents related by *Background* describe eventualities that temporally overlap; there's also an axiom which states they have a distinct common topic.

These conditionals are axioms of the glue logic DICE. As we assume that interpreters use this glue logic, we assume that these axioms are also part of the interpreter's stock of information, together with information given by the transfer function about the semantic content of the discourse context and the current clause, and about his beliefs about other participants in the discourse. We will call this stock of information expressed in the glue language and used by DICE the interpreter's knowledge base (KB), keeping to our earlier usage (e.g. Lascarides and Asher 1993). This usage, we now realize, is misleading; the KB is not to be identified with the interpreter's cognitive state, because the interpreter's beliefs are plausibly expressed in a much richer language than the glue language. To avoid misunderstanding, we will take KB as a technical term—referring just to the stock of information conveyed by the transfer function into the glue language.[6]

DICE exploits a nonmonotonic notion of validity ($\mid\!\approx$), so as to use these default rules to compute the rhetorical relation between constituents that results from the interpreter's KB. This particular notion of validity has several nice properties, which we have discussed in detail elsewhere (Lascarides and Asher 1993). Four of these are relevant for our purposes. First, $\mid\!\approx$ validates Defeasible Modus Ponens (DMP): when the default laws whose antecedents are verified in the KB all have consequents that are consistent with the KB and with each other, then all the consequents are nonmonotonically inferred. Second, $\mid\!\approx$ validates Specificity: when conflicting default rules apply, the consequent of the most specific default rule (if there is one) is inferred. Third, $\mid\!\approx$ is robust in that if $\Gamma\mid\!\approx\phi$ then ϕ will survive as a consequence of the premises Γ augmented with logically independent information. And finally, for each deduction $P\mid\!\approx Q$ there is a corresponding embedded default in the object language (that is, a formula in which one occurrence of the connective $>$ occurs within the scope of another) which links boolean combinations of the formulae P and Q, and which is verified to be true. This amounts to a weak deduction theorem, but we are assured of such embedded defaults only for quantifier free languages. It turns out that for analyzing the interaction of cognitive states and discourse structure in dialogue, we need quantification. To be able to express $\mid\!\approx$ within the object language at least partially, we will add to the underlying logic a rule that enables us to derive new defaults from the theory. We will call it the `Weak Deduction Rule`:

- `Weak Deduction Rule:`
  Let T be a theory in L_{DICE}, and suppose $T, P \mid\!\approx Q$ but $T \not\mid\!\approx Q$ and $T \not\models \neg(P > Q)$.
  Then $T \mid\!\approx P > Q$.

`Weak Deduction` preserves consistency but results in a nonconservative extension of DICE, as discussed in Lascarides and Asher (1993). We will use `Weak Deduction` in section 7 to deduce default axioms of DICE from more general principles about how beliefs, desires and intentions interact, which we present in section 4.2.

To see how DICE works to compute rhetorical relations, consider (9). The compositional semantics of (9) includes the information that the main eventuality described by the clause (9b) is a state (cf. Kamp and Reyle, 1993). Therefore, via the transfer function μ, we know that the antecedent to `Background` is verified in the interpreter's KB. Moreover, its consequent is consistent with that

[6] As a general policy, Gricean maxims are either derivable from the axioms in DICE, or epiphenomenal (for details, see Lascarides and Asher, 1993). In fact, we don't take any previous theory of conversational principles such as Grice's or Searle's for granted.

KB. So DMP on `Background` and then Modus Ponens on `Axiom on Background` will correctly predict that (9b) provides background information for (9a) and temporally overlaps with it.

(9) a. Max opened the door.

b. The room was dark.

On the other hand in (10), even though the antecedent to `Background` is verified, one can't attach the constituents with *Background* because the constraint that there must be a distinct common topic would be violated, and hence the consequent of `Background` is inconsistent with the KB. Furthermore, the other rules in DICE for computing rhetorical relations fail to supply an alternative, for either the antecedent of the relevant default rule isn't verified, or coherence constraints on the rhetorical relation are violated. So a rhetorical relation cannot be inferred, leading to incoherence.

(10) a. Max opened the door.

b. ?Mary's hair was black.

In monologue, a discourse is defined to be *incoherent* if one cannot compute an SDRS for that discourse from the interpreter's KB in DICE. Thus our notion of discourse incoherence is dependent on the interpreter's background knowledge, as well as the compositional semantics of the discourse (since this forms part of his knowledge). Discourse coherence is essentially an interpreter-relative property. We'll see in section 4.2 that in dialogue, the conditions under which a discourse is coherent are extended slightly.

SDRT also pays particular attention to the way discourse structure affects the content of a discourse, understood in a dynamic fashion. A distinctive feature of SDRT is that the rhetorical relations constrain the semantic content of the constituents they connect together, and these constraints can trigger *modifications* to the content of the constituents as discourse is processed. If information that is used to establish the coherence of a rhetorical relation is not already there, then it's added in a constrained manner. For example, we inferred the additional semantic content that the event and state overlap in (9), which wasn't given by the compositional semantics alone.[7] Intuitively, these modifications record the *accommodation* that's necessary to interpret utterances correctly. Asher (1993) shows how encoding these modifications explicitly is useful in assessing how subsequent portions of discourse relate to the previous ones. Our analysis of (1) will involve these tools for accommodation.

Having completed this quick introduction to the glue logic DICE, we can now return to the way information is updated in SDRT, and compare it with the *Update* function in DRT. Similarly to DRT, $Update_{sdrt}$ is defined with respect to an SDRS K which represents the given context, and an SDRS K' which represents the new information to be incorporated into it. But unlike DRT, updating old information with new in SDRT is defined as a *relation* rather than a function. One reason is that there may be several available attachment points in the old information K to which the new information K' could attach with a rhetorical relation. As discussed in Asher (1993) and elsewhere, the available attachment points in K are those on the right frontier, where the subordinating relations are *Elaboration*, *Explanation* and $\Downarrow$ ($\alpha \Downarrow \beta$ means that α is a topic for β). So there is some indeterminacy in computing the update of K with K'.

Because of this, we define an $Update_{sdrt}$ relation, which relates K, K' and a new SDRS $K^\dagger$; from an intuitive perspective, $K^\dagger$ is an SDRS where the old and new information have been merged together. More specifically, $K^\dagger$ includes (a) the old information K, (b) the new information

[7] This contrasts with the analysis of this discourse given in DRT (Kamp and Reyle, 1993), where the overlapping relation is part of the compositional semantics of this discourse. But we have argued extensively elsewhere that this approach is problematic for other discourses involving events and states (e.g., Lascarides and Asher, 1993).

K', and (c) an attachment of K' with a rhetorical relation to an available attachment point in K, where (d) K and K' now include any semantic content that we add in order to satisfy that rhetorical relation's coherence constraints (e.g., in the case where we attach the constituents with *Background*, the temporal overlap between the eventualities would be added if it's not already there). The relation $Update_{sdrt}(K, K', K^{\dagger})$ is constrained so that it can hold only if parts (c) and (d) of $K^{\dagger}$ are computed via the axioms in DICE. So axioms like `Background` serve to constrain the way an SDRS gets updated with new information.

More formally and following Asher (1996), the relation $Update_{sdrt}(K, K', K^{\dagger})$ is defined as follows. Let $Avl(K)$ be the set of available attachment sites (which are labels π) in the SDRS K. And let $\mid\!\sim$ be the proof theoretic counterpart to the DICE consequence relation. 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, let K_{π} stand for the SDRS constituent labelled by π, and finally let $\alpha[\beta/\gamma]$ be the result of replacing γ in α with β. Then the $Update_{sdrt}$ function is defined as:

- `The Update Relation`
  $Update_{sdrt}(K, K', K^{\dagger})$ is true iff $\exists\pi \in Avl(K)$ such that:
  1. $(\mu(K_{\pi}), \mu(K'))\mid\!\sim(R(\pi, \pi') \wedge \varphi)$; and
  2. $K^{\dagger} = K_{\tau}[K^{+}/K_{Pred(\alpha)}]$, where:
  3. $K^{+} = Update_{drt}(K_{Pred(\alpha)}, \langle\{\beta\}\{\beta : K_{\beta}(\varphi), R(\alpha, \beta)\}\rangle)$, where $K_{\beta}(\varphi) = K_{\beta}$ together with those conditions specified in φ, where φ is that information needed to satisfy the coherence constraints on R.

So the resulting SDRS for the discourse will incorporate a revision of the new information and this newly revised bit of information will be attached to some attachment point via a rhetorical relation R that's computed on the basis of the axioms of DICE together with the knowledge of the semantic content of the old and new information (i.e., $\mu(K_{\pi})$ and $\mu(K')$), where the SDRS labels π and π' stand for the SDRSs K_{π} and K', once these SDRSs have been modified by adding the semantic content φ that was necessary for ensuring that the coherence constraints on R are met.

Asher (1996, forthcoming) shows that SDRS updates so defined correspond to transitions on "more model-theoretic" conceptions of discourse contexts than SDRSs. This replicates the correspondence between DRS update and the transition function on information states as sets of model embedding function pairs proved in Fernando (1994). However, updates in SDRT are more complex than in DRT; there is more than set union going on. First, there is indeterminacy on where new information attaches in the old SDRS. Second, the way it attaches is determined by axioms in DICE; and DICE is designed to model the way semantic content and background knowledge affect the connections between propositions, and the way those connections affect the content of the propositions themselves. The modifications that we will make to DICE in this paper will make it possible to prove that the equivalence of transitions on information states is no longer r.e., in the way that it is proved to be in Fernando (1994) for DRT and Asher (1996) for various versions of SDRT in monologue. Essentially what we will do below is to introduce non-eliminable quantification into the axioms of DICE, and this will render our SDRS update more difficult to compute.

4.2 Extending SDRT to Dialogue

We wish to use SDRT to model dialogue, so as to take advantage of the way it models the information flow between compositional semantics and discourse structure. To see how we might use SDRT, let's go back to our original example:

(1) a. A: How can I get to the treasure?

b. B: It's at the secret valley.

c. A: But I don't know how to get there.

Attaching (1b) to the discourse context (1a) with a rhetorical relation will trigger modifications to (1b) and the cognitive state of the interpreter via formally precise reasoning. Roughly, (1b) in this context will be interpreted as: *The treasure is at the secret valley, you can get to the treasure by going to the secret valley, and I assume that you know how to go to the secret valley.* These modifications will then affect how the representation of (1c) attaches, because the coherence constraints on the rhetorical relation *Contrast*, as signalled by the cue word *but*, are dependent on the way (1b) is interpreted.

To model dialogue, we will have to complicate the picture. In general, the participants A and B of a dialogue have different cognitive states. And since the SDRS is determined using the glue logic DICE which exploits information in the interpreter's cognitive state, the SDRSs they construct in general differ. So as discourse analysts, we must build at least *two* SDRSs; one for A and one for B. In fact, A and B must also do this, since they each have models of each other's cognitive states. When there is no miscommunication or misunderstanding, A's and B's SDRSs are isomorphic. When misunderstandings occur between interpreters who understand the words in the dialogue, the two SDRSs are not isomorphic. Dialogue participants detect such misunderstandings when their own SDRS does not match the SDRS they construct from their model of the other participant's cognitive state.[8] They then revise both their beliefs about the other participants as well as their representation of the discourse context. As we extend SDRT to analyse dialogue, we'll extend the tools for SDRS accommodation so that they can be used to modify the cognitive states of the participants, as well as the semantic content of the constituents.

It is important to keep track of who said what in dialogue, if we are to get an accurate picture of the interaction between cognitive states and discourse structure. So each agent tags the constituents in his SDRS with who said it: $A{:}\alpha$ means that participant A has said that α. We will write α for $A{:}\alpha$, when there's no confusion. And to simplify notation, we'll ignore for the rest of the paper the labels of constituents that we discussed in the last section.

By emphasizing the interactions between the participants' cognitive states and discourse structure in dialogue, a new perspective on rhetorical relations emerges. As in monologue, it is natural to think of rhetorical relations as linking instances of structured propositions. But in addition, we naturally think of these instances as forming a pattern of responses; each speech act by an agent is linked to another speech act (perhaps of a different agent) in one of many intricate ways.[9] While the question-answer pattern is one central pattern in dialogue, there are also others—corrections of perceived misunderstandings, comments on other responses, etc.

This perspective makes each relation correspond to a *speech act type*, where the second term of the relation is a speech act of the appropriate type relative to its discourse context, which of course includes the first term of the relation. For example, *Explanation* is a particular type of speech act, and in $Explanation(\alpha, \beta)$, β is the speech act, performed by the person who uttered β, of asserting an *Explanation*, relative to the discourse context α. When extended to dialogue, SDRT will enrich and complement extant theories of speech acts by linking individual speech acts within a larger discourse structure and systematically building links between that structure and the agents's cognitive states. Whereas traditional speech act theory is largely sentence bound, SDRT would serve a theoretical basis for investigating speech acts and their effects at the suprasentential

[8] We remain agnostic as to whether a common ground SDRS, following the lines of (Stalnaker, 1978) and (Cohen and Levesque 1990, Lochbaum 1993, Green and Carberry 1994, and others) should be built as well.

[9] We could also think of the discourse structure as a discourse plan using the terminology from from AI (e.g., Perrault 1990). But though the planning terminology is suggestive, little in the way of inference peculiar to planning rather than standard nondynamic logic seems to be applied in extant theories of dialogue (e.g., Perrault 1990, Grosz and Sidner 1990, Lochbaum 1995, Pollack 1990, Litman and Allen, 1990).

(or discourse) level.

Viewing relations in SDRT as linking speech acts emphasises that we must define their semantics at least in part in terms of the cognitive states of the participants. Viewing an SDRS this way also makes the SDRT approach to dialogue more compatible with Grosz and Sidner's, but with SDRT paying more attention to the interaction between lexical and compositional semantic content and discourse structure. The gaping hole in SDRT's analysis of dialogue is an examination of the interaction between this semantic information and mental states. Our aim is to do this for the question-answer alternation.

The interaction between an agent's cognitive states and discourse structure, as we've pointed out, will make a difference to SDRS construction; and so we must enable the glue logic DICE to use information from agent's cognitive states. To this end, we now sketch five additions to DICE.

First, we add to DICE attitude operators for belief, desire and intention and three rules about belief, which enable A and B to use the axioms in DICE to believe that a discourse has the discourse structure they have computed, and to come to conclusions about the nature of the other participants' SDRSs. We shall show how they do this in detail in section 7. The first of these belief axioms states that agents believe not only all the monotonic consequences of their beliefs but the nonmonotonic consequences as well. The axiom `Believing Inferences` ensures this (KB_A stands for A's KB; and the semantics of $\mathcal{B}_A\phi$ is that of a simple K4 style operator):

- `Believing Inferences`: $\mathcal{B}_A\phi$ *if and only if* $KB_A \mid\approx \phi$

The second two of these three axioms are about *mutual belief*. As we said above, agents construct SDRSs both for themselves and for the other participants. The participants in a discourse mutually believe that they are all doing this. So we will assume that the set of agents A, B in the dialogue mutually believe the axioms of DICE; `Mutually Believing DICE` stipulates this. Finally, we formalize mutual belief as in Dwork and Moses (1986). `Mutual Belief` states that if A and B mutually believe ϕ, then A believes ϕ and also believes they mutually believe it, and similarly for B.

- `Mutually Believing DICE`: $MB_{A,B}\text{DICE}$
- `Mutual Belief`: $MB_{A,B}\phi \rightarrow (\mathcal{B}_A(\phi \wedge MB_{A,B}\phi) \wedge \mathcal{B}_B(\phi \wedge MB_{A,B}\phi))$

`Mutual Belief` is an axiom which guarantees that $MB_{A,B}\phi$ entails $\mathcal{B}_A\phi$, $\mathcal{B}_A\mathcal{B}_B\phi$, $\mathcal{B}_A\mathcal{B}_B\mathcal{B}_A\phi$, etc *ad infinitum*.

The second addition to DICE is constraints about how the different agents' SDRSs in dialogue cohere together. As we've mentioned, SDRT for monologue predicts that a discourse is coherent for an interpreter if he can compute an SDRS for it, via the rules in DICE. This holds for dialogue too. But in addition, an interpreter is not only constructing his own SDRS, but also the SDRSs for the other participants from his model of their KBs, so that he can reason about what they think the dialogue means. Thus the definition of coherence for dialogue is more complex: the current utterance is coherent for an interpreter if he can compute not only a rhetorical relation that connects it to his SDRS, but also a rhetorical relation that he thinks the speaker intended. This latter condition captures the intuition that the utterance is coherent for the hearer only if he can compute what he thinks the speaker intended to convey by it. `Computability`—which is a DICE axiom on what are possible SDRT attachment sites for β—captures this.

- `Computability`:
  A can add $\langle\tau, \alpha, B{:}\beta\rangle$ *to his* KB KB_A *only if for some* $R \in$ *rhet-relns*,
  $KB_A, \langle\tau, \alpha, B{:}\beta\rangle \mid\approx \mathcal{I}_B\mathcal{B}_A(R(\alpha, \beta))$.

In words, `Computability` states that the hearer A can assume B's utterance is to be attached to α only if he has sufficient information in his KB KB_A to come to a default conclusion about what the speaker B intended A to believe was the rhetorical connection between his utterance and α.[10] That is, the hearer must have a default idea about what the speaker was trying to contribute to the dialogue. Intuitively, this corresponds to the hearer's assumption that the speaker is cooperative.[11] We'll see that `Computability` is violated when attaching (1b′) to (1a), leading to incoherence. But in contrast, it's satisfied for (1b).

Third, we need to add some formal means for reasoning from the cognitive states of participants to what they say and back again. In Asher and Lascarides (1994), we used a simple scheme which relates desires, beliefs and intentions, and which is known as Aristotle's Practical Syllogism, to model the information flow between semantic content and cognitive states in simple monologues.[12] We use The Practical Syllogism here to model the information flow in dialogue. Aristotle's Practical Syllogism states that normally, people intend to do things that they believe help them achieve their goals. The DICE axiom below describes this. If (a) B wants ψ and believes $\neg\psi$, and (b) B believes that he can nonmonotonically infer ψ from his KB augmented with ϕ, and moreover ϕ is B's choice for achieving ψ, then normally (c) B intends ϕ.

- `Practical Syllogism:`
  - (a) $(\mathcal{W}_B(\psi) \wedge \mathcal{B}_B(\neg\psi)\wedge$
  - (b) $\mathcal{B}_B((\phi > \mathit{eventually}(\psi))) \wedge \mathit{choice}_B(\phi,\psi)) >$
  - (c) $\mathcal{I}_B(\phi)$

Finally, we are going to restrict our investigation to dialogues where both participants are being cooperative, and in particular that normally, they utter only what they believe to be true (cf. Grice's Maxim of Quality). We therefore assume axioms of cooperativity and sincerity, given below:

- `Cooperation:` $(\mathcal{W}_A p \wedge \mathcal{B}_A \neg p) > (\mathcal{W}_B p \wedge \mathcal{B}_B \neg p)$
- `Sincerity:` $\mathcal{I}_B\mathcal{B}_A(R(\alpha, B:\beta) > \mathcal{B}_B(R(\alpha, B:\beta))$

In words, `Cooperation` states that B normally adopts A's goals (cf. the constraints on goalhood given in Cohen and Levesque, 1990). `Sincerity` states that the speaker B intends to make A believe a connection R between his utterance and the context normally because he believes this connection himself.[13] Note that `Computability` and `Sincerity` entail the following:

- *A can add $\langle\tau,\alpha,B{:}\beta\rangle$ to his* KB KB_A, `Sincerity` $\approx\!\!\!\mid\; \mathcal{B}_B(R(\alpha,\beta))$.

These axioms and `The Practical Syllogism` play an essential part in the reasoning behind how responses to questions are interpreted in dialogue. In section 7, we will use them to show that, together with our account of how the compositional semantics of interrogatives and answers affect discourse structure, we can derive the principle that in general a respondent thinks he's answering questions appropriately for the questioner. That is, he thinks his response will enable the questioner to achieve the goals that lay behind the question. In this sense, we derive in a

[10] This is a very weak assumption, given that one can take some of the rhetorical relations in *rhet-relns* to be very vague or general semantically, and that the above rule doesn't block A from being able to infer several relations between α and β; he simply has to infer at least one.

[11] We ignore the analysis of dialogues in non-cooperative contexts in this paper.

[12] This scheme abstracts away from many details and won't work in cases where we have conflicting desires. But it illustrates the sort of interactive principle that one could use. A more thoroughgoing approach might use full-blown decision theory.

[13] Obviously, both these axioms would need to be dropped if we were to analyse uncooperative dialogues, such as dialogues where people deliberately lie, mislead or change the course of the conversation in order to block a participant from achieving his goals. This is a matter for future investigation.

principled fashion the goals and cognitive reasoning that is required in a semantic account of interrogatives like Ginzburg's (1995), thereby filling a gap in his theory. This is discussed in detail shortly.

5 The Semantics of Interrogatives and Answers

Since in building discourse structure we wish to exploit monotonic information about compositional semantics whenever it is possible to do so, we require a compositional semantic analysis of interrogatives and answers. As we mentioned, we will concentrate on *how*-interrogatives. Extant formal semantic theories have largely ignored these. We propose an analysis of *how*-interrogatives and their answers that builds on the previous formal semantic work on interrogatives.

An elegant picture of the formal semantics of interrogatives comes from Belnap (1985), Bennett (1979), Karttunen (1977) and Groenendijk and Stokhof (1982, 1984). For Karttunen and Bennett, the meaning of an interrogative is a function $[\![q]\!]$ from worlds to sets of propositions, where propositions are sets of worlds. For Karttunen, the conjunction of all elements in $[\![q]\!](w)$ is glossed as the true, (weakly) exhaustive answer at w, by which we mean, for instance, that the meaning of *who came to the party?* must specify all those individuals who came to the party. It is *weakly* exhaustive in the sense that although an answer must specify all those individuals who came to the party, it need not specify all those individuals who did not. We follow this intuition and will speak of propositions in the denotation of a interrogative as being answers to the question.

Groenendijk and Stokhof use λ abstracts over indices to represent interrogatives. *wh*-interrogatives have a particular form that involves a further λ abstract, which serves a rather different purpose; they bind a variable that is the translation of the trace of the *wh*-phrase. The kind of variable bound will depend on the type of interrogative. For instance, the interrogative *Who came to the party?* binds a variable of individual type. The meaning of this interrogative at index i is the function from indices to the set of indices A such that for each $a \in A$ $[\![\lambda x loves'_i(x, m) = \lambda x loves'_a(x, m)]\!]$. So for instance if John and Mary and no one else came to the party, then the meaning of the interrogative in such a situation is just the set of indices (the proposition) in which John and Mary and *no one else* came to the party. Notice that this (strong) form of exhaustiveness entails Karttunen's notion of exhaustiveness.

We will follow Groenendijk and Stokhof and use lambda abstracts to represent interrogatives, but our compositional semantics will be different. As Hintikka (1983) and Ginzburg (1995) point out, the analyses in terms of strongly and weakly exhaustive answers poses problems. Groenendijk and Stokhof (1984) argue for a strong exhaustive analysis of answers on the basis of the truth conditions of sentences containing *knowing who*. Ginzburg (1995) points out that the semantics of other embedded *wh*-interrogatives such as *know where*, *know what* and *know which* require a non-exhaustive analysis of answerhood; what Groenendijk and Stokhof would define as *partial* answers.

It is natural enough to argue that *Mary* contributes a partial answer to the question *Who came to the party?* under those circumstances where Mary and perhaps some others came; and pragmatics may relax the constraint of exhaustiveness for answers. But in some cases, in particular with *how*-questions, as we will argue below, exhaustiveness hardly ever appears to be a requirement on the meaning of the interrogative. In agreement with Ginzburg (1995), we think that this notion of "relaxation" will be quite difficult to capture unless we take the semantics of interrogatives, viz. the *exhaustiveness* requirement on answers, to be some sort of default. But this last assumption seems to get the pragmatics/semantics interaction backward. Ideally, semantics should give monotonic information about content, while pragmatics *supplements* this with defeasible content. An analysis where pragmatics removes exhaustiveness in certain contexts from the compositional semantics of the answer will be one where the compositional semantics potentially provides *more* information than the pragmatic interpretation. It seems rather more likely that exhaustiveness is a pragmatic effect of various Gricean maxims. This idea isn't new—cf. Hintikka (1975) and others. But of course we are required to give a precise account of how such a semantics for answers can be spelled out, and how in spite of the non-exhaustive answers, one still apparently has an exhaustive reading of certain *wh*-questions in embedded contexts.

Like Groenendijk and Stokhof, Karttunen and others, we give a compositional semantics of interrogatives in terms of answers. But we allow propositions that in our terms constitute non-exhaustive answers to the question to play a direct role in specifying the semantics of an interrogative. The set of propositions $[\![q]\!](w)$ corresponding to the meaning of an interrogative q at a world w is the set of all true direct answers to the question, including non-exhaustive ones, that either determine some positive extension of the λ abstract or if it is empty specify that it is empty. Thus, if Mary came to the party, *Mary* is an acceptable answer and part of the meaning of *who came to the party?*, even if other people came.

We will suppose that the logical form of a *wh*-interrogative is a λ abstract where the bound variable occupies the trace position left by the *wh*-element. While we could also have as in Groenendijk and Stokhof a λ abstract over indices or over propositions, we'll remain agnostic about incorporating such types into our logic. Instead, the logical form of *Who loves Mary?* has the logical form $?\lambda x loves'(x, mary')$. The set of direct answers associated with this question at a world w are all of the form $\alpha(\lambda x loves'(x, mary'))$, such that $w \models \alpha(\lambda x loves'(x, mary'))$, where α is any expression of the type that takes a property as an argument, and returns a proposition (typically a generalized quantifier denoted by an NP). We add one constraint—namely that the answer must either entail that $[\![loves'(x, mary')]\!]$ is nonempty at the world of evaluation or if it is empty the answer must entail that. Moreover, we agree with the claim made at least since Hintikka that these propositions must entail that some particular objects occur in the extension of *loves Mary*, thus supporting a *de re* knowledge claim when the question is embedded in attitudinal contexts.[14]

More formally, *?* is a function from dynamic n-ary relations $\lambda x_1, \cdots, \lambda x_n \psi(x_1, \cdots, x_n)$ (where dynamic properties and relations are λ-abstracted SDRSs, as described in Asher (1993)), to a function from a world w into a set of propositions as defined below in (11):

(11) $$\{p : p = [\![\alpha_1, \cdots \alpha_n(\lambda x_1, \cdots, \lambda x_n \psi(x_1, \cdots, x_n))]\!] \wedge^{\vee} p \wedge (\exists x_1, \cdots x_n \Box(^{\vee}p \rightarrow \psi(x_1, \cdots, x_n)) \vee \Box(^{\vee}p \rightarrow \neg\exists x_1, \cdots x_n \psi(x_1, \cdots, x_n)))\}$$

That is, each proposition p in the set fills the λ-abstracted arguments, is true, and either entails that some particular objects occur in the extension of the λ-abstract, or that no objects occur in the extension. For example, a simple *yes-no* interrogative $?\phi$ will have as a meaning a function from worlds to a singleton set of propositions—for a given world w the set will contain the proposition ϕ if $[\![\phi]\!] = 1$ and $\neg\phi$ otherwise. For a *wh*-question like *who loves Mary?* (ignoring tense for now), the set becomes that given in (12):

(12) $$[\![?\lambda x loves(x, m)]\!](w) = \{p : p = [\![\alpha(\lambda x loves(x, m)]\!] \wedge^{\vee} p \wedge (\exists x \Box(^{\vee}p \rightarrow loves(x, m)) \vee \Box(^{\vee}p \rightarrow \neg\exists x loves(x, m)))\}$$

That is, the answers to this interrogative are: the set of propositions p such that for each p, either p entails that some particular person came, or for each p, p entails no one came. For us, these propositions will be SDRSs. Thus *John loves Mary* is an answer to *Who loves Mary?*, so long as *John loves Mary* is true.

However, *Not Bill* is not a direct answer to *Who loves Mary?* because *Not Bill* does not entail either that someone loves Mary, nor that no one loves Mary. And in contrast with Ginzburg,

[14]The inference formally is this. Suppose knowledge is closed under entailment and that $K\phi$ and $\exists x \Box(\phi \rightarrow \psi(x))$. Then it follows in most standard epistemic logics that $\exists x K\psi(x)$. In the DRT semantics for attitudes, this will not follow without further assumptions. We believe that in the case of these entailments, these assumptions are justified, but we will not go into the matter here.

our definition does not allow the DRS representing *Several politicians love Mary* to be an element of $[\![?\lambda x(loves'(x, m'))]\!]$, because this response doesn't specify a particular extension to the set of lovers of Mary. Specifying a particular extension is important for answers to questions: it is very important in communication in dialogue, as well as in interrogatives embedded within attitude contexts, as we will see in the next section.

In light of this semantics, it's also important to stress that analysing questions this way doesn't preclude false answers from forming coherent responses in dialogue in our analysis. We will see in section 6 that our analysis of dialogue predicts, correctly, that the questioner A in (13) finds B's response coherent, even if he knows that this response isn't true:

(13) A: Who does Bill like?
B: Mary.

So true answers aren't a necessary requirement for coherence, in spite of the role that true answers play in the semantics of interrogatives. The two concepts truth and coherence are obviously related, however. We'll spell out the role truth plays in determining the structure of the dialogue in section 6. Briefly, the rhetorical relation A uses to attach B's response to his question will be affected by whether he has reason to believe the response is false. And this difference in rhetorical structure can serve to influence what A should say next. At any rate, the fact that A will be able to attach B's response (which A believes to be false)to his interrogative with a rhetorical relation in his SDRS, will mean that this dialogue is coherent for A. Details about exactly which rhetorical relation is used in these circumstances are given in section 6.

Sometimes an adequate response to an interrogative in discourse may not constitute a direct answer in our sense. In the dialogues (14) and (15), the responses might be appropriate, in the sense that they give enough semantic content that A can now fulfill the goals that lay behind his question, though they aren't direct answers, given the compositional semantics of direct answers that we've provided.

(14) a. A: Who came to the party?
b. B: Not John.

(15) a. A: How do I get Downtown?
b. B: Well, you get there either by taking I-79 or I-279, but I'm not sure which.

(14b)/(15b) aren't direct answers because they don't entail that someone came or no one came/that there is some particular adverbial of manner for getting Downtown, or no adverbial of manner. They are, however, *partial answers*, which we define in the spirit of Groenendijk and Stokhof: they reduce the space from which one can find direct answers. They therefore help the questioner by giving him some information about the subject he wanted to know something about. As such, we will ensure that they are coherent responses; that is, the interpreter will be able to attach the response to the interrogative with a rhetorical relation. We return to this discussion in section 6.2.

It should be stressed that partial answers like those in (14b) and (15b) are distinct in our theory from indirect answers. For us, indirect answerhood is a context sensitive notion, dependent on the knowledge base of the questioner. In section 6.2, we will formally define indirect answers to be those responses which, when added to the questioner's knowledge base, permit the questioner to infer a direct answer from it. It's not necessarily the case that A can infer a direct answer to question (14a) from the content given in (14b) and everything else he knows. If he can't, then (14b) is a partial answer, but not an indirect one. If he can, then it's both a partial answer and an indirect one.

We hope that this discussion clarifies the similarities and differences between our semantics of interrogatives and that of Groenendijk and Stokhof. We make now a similar comparison between our work and that of Ginzburg (1995). Ginzburg uses his arguments against the type of analysis of interrogatives provided by Groenendijk and Stokhof to motivate a radically new strategy for defining the compositional semantics of interrogatives. It is radical, in that he doesn't supply a semantics based on a notion of answerhood that is exploited compositionally. Rather, he argues that one needs several notions of answerhood to represent the semantics of interrogatives. Among these is the notion of being information that fully resolves a question q. This is needed for the semantics of interrogatives embedded by resolutive predicates; for instance *know*, in *Max knows who came to the party*. Fully resolving a question is defined in terms of the interpreter's goal and mental state. It is designed to capture the fact that in the context of (16a,c,d), one can truthfully say (16e), but in the context of (16b,c,d), (16e) is false.

(16)
a. [Context: A has just travelled by air]
b. [Context: A has just taken a five minute taxi ride]
c. A: Where am I?
d. B: Helsinki.
e. B knows where he is.

Since *fully resolves* is defined in terms of the questioner's mental state and goals, it is an extremely context sensitive notion. Another notion of answerhood that Ginzburg defines is being information that *potentially* resolves a question q. This defines the information that could resolve the question. In essence: there is some goal which prompted the question, and some mental state that the interpreter could have, such that this information would resolve the question for him. Thus in the above example, *Helsinki* potentially resolves the question, even if it doesn't fully resolve it. Thus potentially resolves abstracts away from the particular mental state and goals at hand in the particular context in which the interrogative is uttered.

Our notion of answerhood is similar to his notion of potential resolvedness. To capture the data Ginzburg describes concerning the appropriacy of responses and the satisfiability of embedded interrogatives in different contexts, we will extend our analysis of questions by embedding our semantics in SDRT. For instance to explain the context sensitivity of responses that Ginzburg observes, we will exploit the SDRT representation of dialogue contexts, which includes the representation of the cognitive states of the participants. Ginzburg's context-sensitive interpreter-specific notions of answerhood are not defined directly in our analysis, but they're reconstructed via a combination of SDRSs and cognitive states, as we will see shortly. These more pragmatically rich representations of the meaning of responses can be used as arguments in propositional attitudes, so we in principle could explain the data concerning embedded interrogatives that Ginzburg discusses.

Our analysis also takes Ginzburg's story one step further. Ginzburg doesn't define what goals are; nor does he link them to plans, beliefs, desires and intentions. He also doesn't explain how what's said in a dialogue can contribute to inferences about what the hearer thinks the speaker thinks. Therefore, he doesn't model what the hearer thinks the speaker thinks about the hearer, making it difficult to detect misunderstandings on the basis of what's said. Ginzburg also doesn't explain how responses to questions can help one revise plans. Consequently, he can't fully explain dialogue (1), because he doesn't link the response (1b) to the action of going to the secret valley, via complex spatial reasoning.[15] Finally, Ginzburg doesn't account for the interaction between questions and anaphora in dialogue. Our approach will fill in these gaps in his analysis.

Further, we will derive the context sensitive notions of when a response resolves a question relative to the questioner's needs etc. from the compositional semantics of answers, plus general principles about how people think, which were given in section 4.2—we will call such a response *appropriate*.[16]

[15] Groenendijk and Stokhof (1984) also don't supply a theory of how the participants' cognitive states change as the dialogue progresses.

[16] We do not follow Ginzburg (1995) in using situation theory, for several reasons. First, we want to use SDRT in

By doing this, we aim to add to the insights provided by Ginzburg (1995), in that our theory will use very general axioms about beliefs, desires and intentions that were given in section 4.2 to track how the participants' cognitive states change as the dialogue proceeds. Thus, we aim to preserve some compositional semantic notion of answerhood, and yet still explain, in terms of the context sensitive notions of goals and mental states, the circumstances when a response resolves a question relative to the questioner's needs (or in our terms, when a response is appropriate).

5.1 The Problem of Embedded Interrogatives

Analysing questions in terms of (perhaps) non-exhaustive answers receives support for embedded *how*-interrogatives, for (17) is true so long as John knows one way of getting to the treasure:

(17) John knows how to get to the treasure.

If the semantics of a *how*-interrogative was defined in terms of exhaustive answers, then John would have to know *all* possible ways of getting to the treasure, in order to know how to get there. So the proponents of weak *and* strong exhaustiveness in the semantics of interrogatives would appear to make incorrect predictions for such embedded interrogatives. We (as well as Ginzburg) don't have a problem with the non-exhaustivness in the natural interpretation of sentences like (17). Nor do we have problems interpretating sentences such as (18), where intuitively, Jane does not need to know the most precise characterization of her location (down to quantum indeterminacy?).

(18) Jane knows where she is.

Nevertheless, there is a minimum threshold on the amount of content that John needs to know in order to know how to do something/know where something is. In the context of dialogue (15), for example, even after A has interpreted utterance (15), he generally *doesn't* know how to get Downtown, because (15b) isn't a direct answer. He will know how to get Downtown in the context of (15), however, if he already knows that you can't get Downtown by taking I-79. This is because he will be able to compute a direct answer from (15b); he will know the direct answer *You can get Downtown by taking I-279.*

Under the assumption that *how-*, *where-*, *which-*, *who-* and *why*-interrogatives are all to be uniformly treated as *wh*-interrogatives, this data concerning embedded *how-* and *where*-interrogatives strongly suggests that interrogatives in general must be analysed in terms of non-exhaustive answers.

In contrast to (17) and (18), embedded *who*-interrogatives such as (19) appear to have an exhaustiveness requirement, which we can't capture via our semantics of *who*-interrogatives alone, since their meaning doesn't constitute the set of exhaustive answers.

(19) John knows who came to the party.

We're committed to the non-exhaustive analysis of *who*-interrogatives, because of the data concerning other *wh*-interrogatives. But at first glance, it would appear that (19) isn't true if John

order to exploit constraints on anaphora resolution which are already defined there. Second, we feel that answers to certain *wh*-questions—*why*-questions in particular—must be represented in terms of rhetorical relations. An answer to *?why* α is an explanation, or more formally, a proposition of the form *Explanation*(α, β). In contrast to situation theory, representing such rhetorical information in a DRT environment has been studied extensively (Asher, 1993). And finally, we will exploit the account of propositional attitudes given in DRT, and in particular use the axioms on beliefs, desires and intensions given in section 4.2, to track the cognitive states of the participants as the dialogue progresses. In this way, we hope to enrich some of Ginzburg's insights into the semantics of interrogatives, with general reasoning about anaphora and cognitive states in dialogue.

knows only some of the people that came to the party, but not all of them. In fact, in line with Groenendijk and Stokhof, one might think that (19) isn't true if John believes falsely of someone else that he came to the party. So *knowing who* seems to come with a strong exhaustiveness requirement. Indeed, Gronendijk and Stokhof's strongest argument in favour of a strong exhaustive analysis of *wh*-questions concerns data involving *knowing who*.

There is an account of this strong exhaustive interpretation of *knowing who* to which one can appeal even in the absence of a strong exhaustive analysis of *wh*-interrogatives. We can look at associated pragmatic information. Let us divide the problem into two cases. The first where no one came to the party is easy; both our account and Groenendijk and Stokhof's coincide. It's also easy to get the appropriate knowledge claim from the meaning of the interrogative since the meaning contains just one proposition and intuitively it is that proposition that is known.

The second case is where some people at least came to the party—say Tom, Dick and Harry. In this case, the set of answers (in our sense) form a partial order under entailment, with the topmost element (which entails all the others) being the proposition that Tom, Dick and Harry all came to the party and no-one else did. Dalrymple *et al.* (in press) state an important principle about interpretation. They call it the strongest interpretation principle; it says that you pick the interpretation that logically entails the other. We think that this is right with one proviso—the interpretation must be such that it is compatible in knowledge contexts with the cognitive task at hand. In the case of *who*-interrogatives, the cognitive task is in general compatible with the strongest interpretation of the knowledge claim, as given by the maximally specific proposition in the meaning of the interrogative. This principle thus yields the strongly exhaustive proposition as the one that is known in (19).

But in some cases, particularly (17) and (18), we may know that it isn't reasonable to demand this strongest interpretation, because of the cognitive task at hand. That is, we don't demand that John knows all possible ways of getting to the treasure or that Jane need know the most precise characterization of her location. What we do demand in (17) is that there be some way that John knows and that this way will get him to the secret treasure. In order for Jane to know where she is, we demand a similar *de re* knowledge claim, and this will serve to distinguish the cases given in (16). So much is guaranteed by the semantics of *wh*-interrogatives that we have presented, since in both these cases, (17) and (18) are true only if John knows a direct answer.

Indeed, there are even contexts where the strongest interpretation of an embedded *who*-interrogative isn't compatible with the cognitive task at hand. And so, contrary to what one might first think, there are contexts where embedded *who*-interrogatives lack an exhaustiveness requirement. As an illustrative example, consider sentence (20) in the context where Jill is a gossip columnist, whose task is to write a newspaper article which includes the definitive list of the celebrities that were at the party:[17]

(20) Jill knows who attended the post-Oscar party at Elton John's house.

Then intuitively, (20) is true even if Jill doesn't know that a cameraman from a local TV station attended the party.

To summarise, *John knows ?ϕ* is true only if John knows a direct answer, where these are as defined in section 5, and include non-exhaustive answers. But in addition, there may be more to the demands on particular types of knowledge than this. This added constraint on knowledge is a matter concerning types of knowledge (some of which may be context sensitive), and does not simply arise from the semantics of embedded questions. We leave the matter here, as further investigations would detract us from our main task—the semantics and pragmatics of questions in dialogue.

[17]We would like to thank an anonymous reviewer for pointing out this example to us.

5.2 The Semantics of How-Interrogatives

We distinguish two kinds of *how*-interrogatives: *restricted how*-interrogatives are those that are accompanied by an adverbial of quantity such as *much*, *often*, *small* and *many* (e.g., (24–26)); *unrestricted how*-interrogatives aren't accompanied by such an adverbial (e.g., (21–23)).

(21) A: How can I get to the treasure?
B: By going to the secret valley and digging under the biggest palm tree.

(22) A: How do you steal pie out from under Mom's nose?
B: Quietly.

(23) A: How is Susan?
B: Fine.

(24) A: How much do you love Susan?
B: A lot.

(25) A: How often do you visit Susan?
B: Twice a week.

(26) A: How many children do you have?
B: Two.

In some languages, restricted *how*-interrogatives are lexicalised differently from the unrestricted ones. For instance in French, *how much* would be expressed by *combien*, whereas *how* is *comment*. We take these lexical differences in some languages to justify treating the two types of *how*-interrogative differently.

We take the variables bound by all *how*-interrogatives to be adverbials; but the type of the adverbial will depend on the type of the *how*-interrogative. For unrestricted *how*-interrogatives, the adverbial is one of manner. For restricted ones, it's a modifier of an event that realises the event type. So B's responses in (21–26) are all direct answers to the questions A poses.[18]

It's not always clear where the adverbial attaches in the syntactic structure. Indeed, for unrestricted *how*-interrogatives, the scope seems to vary. For instance, the answer in (22) is naturally understood as a predicate of the event and this would suggest that the variable of manner in the question should be adjoined to the VP. On the other hand, this analysis is much less firmly grounded for examples like (21). Given the treatment of gerunds in Asher (1993), the answer in (21) contains an IP, whose translation yields a fact or possibility, not a predicate of events. Given the bottom up construction procedure for DRT in Asher (1993), it makes more sense to attach the manner adverbial to the IP and to relate the fact expressed by the PP with the fact expressed by the main clause. Accordingly, we will treat the adverb in the question as adjoint to the IP. So the syntactic analysis of (1a) (or equivalently, (21)) is given below.

(1) a. How can I get to the treasure?

[18] Given this semantics of *how*-interrogatives we have provided, we limit attention to *know how* expressions of a certain type: they allow 'propositional' paraphrases, which sets them apart from many other uses of *know how*, such as *John knows how to ride a bike*. The distinction between *knowing how* and *knowing that* is a celebrated one in the philosophical literature (e.g., Hintikka, 1975). And the type of *know how*-expressions dealt with here would be regarded as referring to cases of *knowing that*, as opposed to *knowing how*.

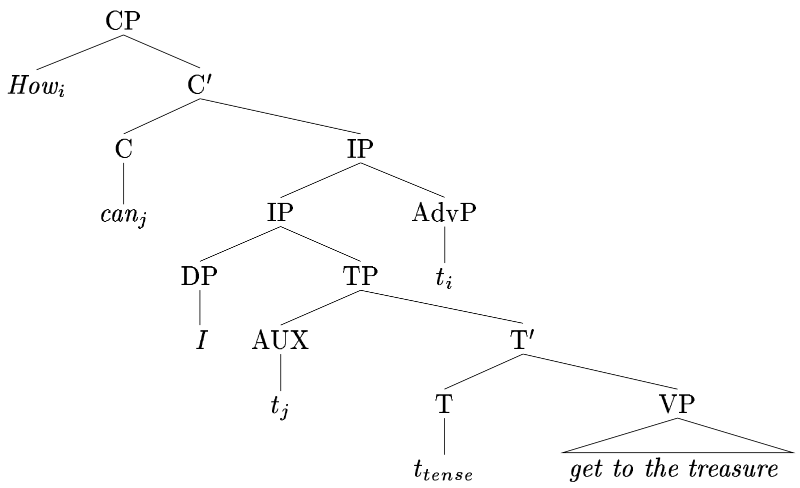

The trace t_j is caused by moving the auxillary can_j to form an interrogative. But we assume, as is standard, that it leaves a variable of AUX type in its trace. This will aid building the compositional semantics, as we will show shortly. The trace t_i would have to be filled by an adverbial of manner expressed as a *by*-phrase containing a gerund phrase of some sort. Filling it with something like *by* ϕ*-ing* produces a direct answer to (1a). A direct answer to (1a) is: I can get to the treasure by going to the secret valley.

In asking (1a), A doesn't require a description of every possible way of getting to the treasure. Nor does our direct answer supply one. Indeed, it's unclear how one could construct the complex adverbials that the semantics of a *how*-interrogative would demand, if one analysed them as referring to exhaustive answers. A direct answer may specify all the possible ways of doing some action—the adverbial of manner may be a complex *by* ϕ*-ing or* ψ*-ing.* But according to our analysis, this isn't necessary. Moreover, as we mentioned before, our analysis captures the intuition that (17) is true if John knows just one way of getting to the treasure.

In building up the DRS for a question, we must decide how the manner adverbial is incorporated into it. In view of our discussion of gerunds and PPs that contain them, our analysis of manner adverbs like that in (21) is distinct from the typical, Davidsonian or neo-Davidsonian account found for instance in Kamp and Reyle (1993). Our hypothesis is that the semantics of the auxillary exploits the contribution of the adverbial. The modal auxillaries *can* and *do* intuitively take the adverbial to be a condition on the realization of the eventuality described by the VP. We realize this intuition technically by interpreting the modal auxillary as connecting a restrictor P which is filled in by a manner adverbial to the nuclear scope Q which is filled in by the tensed VP.

$Can(\phi(x), \psi(x))$ means roughly x *Can* ψ *by* ϕ. Its semantics should capture two intuitions. First, if x actually does ϕ and intends to bring about ψ, he will normally succeed in doing so, eventually. Second, x can ψ by doing ϕ only if x can ϕ. We can motivate this second intuition about *can* via the following simple example:

(27) Alfred can go to the moon by taking a spaceship.

Intuitively, (27) is false if Alfred can't take a spaceship for one reason or another; e.g., because there aren't any spaceships. Ensuring that the semantics of *can* entails that Alfred can do the manner in the restrictor captures this. Formally, we express these two intuitions in the semantics of *Can* below, where a semantics for x's intention that ϕ, written $\mathcal{I}_x(\phi)$, is assumed to be like the DRT-based one given in Singh and Asher (1993).

- `Semantics of Can :`
  $M, w \models_g Can(\phi(x), \psi(x))$ iff
  (a) $M, w \models_g (\phi(e, x) \wedge \mathcal{I}_x(\psi(x))) > \exists e'(e' \succeq e \wedge \psi(e', x))$ and
  (b) $M, w \models_g \mathcal{I}_x(\phi(x)) > \exists e''(\phi(e'', x))$

Clause (a) captures the first intuition and clause (b) the second. Clause (b) does this because by the above semantics of *Can*, $Can(\top, \phi(x))$ (or in words, x can ϕ) is true iff (i) $\mathcal{I}_x(\phi(x)) > \exists e'(e' \succeq e_\top \wedge \phi(e', x))$, and (ii) $\mathcal{I}_x(\top) > \top$.[19] Clause (ii) and $e' \succeq e_\top$ are true trivially. So $Can(\top, \phi(x))$ is true iff $\mathcal{I}_x(\phi(x)) > \exists e''(\phi(e'', x))$ is true, and this is clause (b) in the above definition of $Can(\phi(x), \psi(x))$. Hence x can ψ by ϕ entails x can ϕ.[20]

$Can(\phi(x), \psi(x))$ doesn't entail that $\phi(x)$ or $\psi(x)$ actually hold. On the other hand, the past tense auxillary *Did* seems to have the force of an existential. $Did(\phi(x), \psi(x))$—which should be read as *x did ψ by ϕ*—entails that $\phi(x)$ and $\psi(x)$ both happened in that order in the past and that ϕ was the cause of ψ. We won't give the semantics for *Did*, but it should be obvious (provided one has a semantics for causal connections!) how to do so. We suppose that all auxillaries have either an existential interpretation like *Did* or modal one like *Can*.[21]

The compositional semantic representations of interrogatives are built from the syntax using the "bottom up" DRT construction procedure of Asher (1993). The example (1a) that we are concerned with here is one where the manner adverbial that *how* binds is expressed by a prepositional phrase that includes a gerund phrase. If the manner adverbial is expressed by a subjectless ACC-ing or POSS-ing gerund as in *by going to the secret valley*, then according to the construction procedure in Asher (1993), the gerund phrase will introduce a one-place property with the argument place of this property to be filled in by the controlling subject of the main clause. This property must be an action of the subject, and will be in the restrictor of the modal auxillary *Can*. Assuming that *Can* applies to a tensed VP would imply that the nuclear scope must also be a one-place property of individuals. Thus, the correct logical form of the auxillary *can* should be: $\lambda Q \lambda x \lambda P\, Can(P(x), Q(x))$, where the individual variable x is to be filled in by the partial DRS derived from the subject DP, Q is to be filled by the tensed VP, and P is to be filled by the subjectless gerund phrase.

The DRS construction procedure is shown in Figure 1. We note a few points for clarification. First, we have assumed that the tensed VP *get to the treasure* combines with an AUX variable Δ left by the trace of *can*, but an alternative would be to interpret *can in situ*, prior to movement; the results would be equivalent. This variable Δ is then identified with the semantic contribution of *can*. Second, we have assumed that (1a) is in the present future, which is why the condition $n \prec t$ is in the partial DRS representing the tensed VP. Finally, as is usual in DRT, definite descriptions, indexicals and proper names are promoted to the top level of the DRS, in line with the mechanisms for accommodating presuppositions in DRT which are presented in van der Sandt (1992). The discourse referents i and n stand for the indexicals *I* and *now* respectively. Thus the final representation of (1a) is (1a$'$).

[19] Note that $\top$ is equivalent to $\exists e'' \top(e'', x)$.

[20] Note that the semantics of *Can* does *not* entail that if one does not intend to do ψ and one does ϕ, then ψ won't be achieved. The link from intending to doing something to actually doing it goes only in this direction. And if one does not intend to do ψ, then the semantics of *Can* can still be satisfied. This is necessary for avoiding counterintuitive interpretations of interrogatives such as *How can I calculate* $2+2$?.

[21] One might argue that the above truth conditions of $Can(\phi(x), \psi(x))$ overgenerates. If x can ϕ and x can ψ, then $Can(\phi(x), \psi(x))$ is true as long as x intends to ϕ. This overgeneration results from the fact that the current truth conditions don't record a causal relation between ϕ and ψ; the relation given in clause (a) is somewhat weaker than a causal one. For the sake of simplicity, we have ignored this problem here. The above truth conditions are adequate for the natural language examples that we analyse here.

(1) a. How can I get to the treasure?

a′. $$?\lambda P \begin{array}{|l|}\hline i,n,x \\ \hline \textit{treasure}(x), \\ \textit{Can}\left(\begin{array}{|c|}\hline \\ \hline P(i) \\ \hline \end{array}, \begin{array}{|l|}\hline e,t \\ \hline \textit{get-to}(e,i,x), \\ \textit{holds}(e,t), \\ n \prec t \\ \hline \end{array}\right) \\ \hline \end{array}$$

Similarly, the DRS construction procedure produces a logical form of (28a) that's the same as (1a′) except that $n \prec t$ is replaced with $t \prec n$ (because the tense is past) and the auxillary *Can* is replaced with *Did*.

(28) a. How did I get to the treasure?

P must range over adverbials of manner in both of these. So (1a′) refers to a set of DRSs, where each element of the set replaces P with a one-place property, such as the representation of the subjectles gerund *by going to the secret valley.* (29) is thus a direct answer to the question (1a):

(29) You can get to the treasure by going to the secret valley.

Its syntactic analysis yields the DRS (29′):

(29′) $$\begin{array}{|l|}\hline u,n,x,y \\ \hline \textit{treasure}(x), \\ \textit{secret-valley}(y), \\ \textit{Can}\left(\begin{array}{|l|}\hline e',t' \\ \hline \textit{go-to}(e',u,y), \\ \textit{holds}(e',t'), \\ n \prec t' \\ \hline \end{array}, \begin{array}{|l|}\hline e,t \\ \hline \textit{get-to}(e,u,x), \\ \textit{holds}(e,t), \\ n \prec t \\ \hline \end{array}\right) \\ \hline \end{array}$$

The semantics of *Can* yields several desirable entailments from (29′). First, it entails that you can go to the secret valley. And second, it entails that if you go to the secret valley and you intend to get to the treasure, then normally you will succeed in doing so.

6 From Semantics to Discourse Structure

We now look at how the above compositional semantics of interrogatives and answers contributes to the discourse structure of dialogue. We will do this by embedding this semantics in SDRT, and we'll investigate how it interacts with the discourse coherence constraints already defined there. We'll explicitly mark questions and their direct answers by a new rhetorical relation—the *Question Answer Pair* or *QAP*. We'll then use this relation to define the constraints on responses that aren't direct answers, an example of which is (1b). *QAP* is a rhetorical relation that brings in very directly considerations about why participants ask questions and pragmatic constraints about how to respond to them. These pragmatic considerations will not only be important to the construction of SDRSs but also to understanding the information flow between discourse structure and models of the participants' cognitive states.

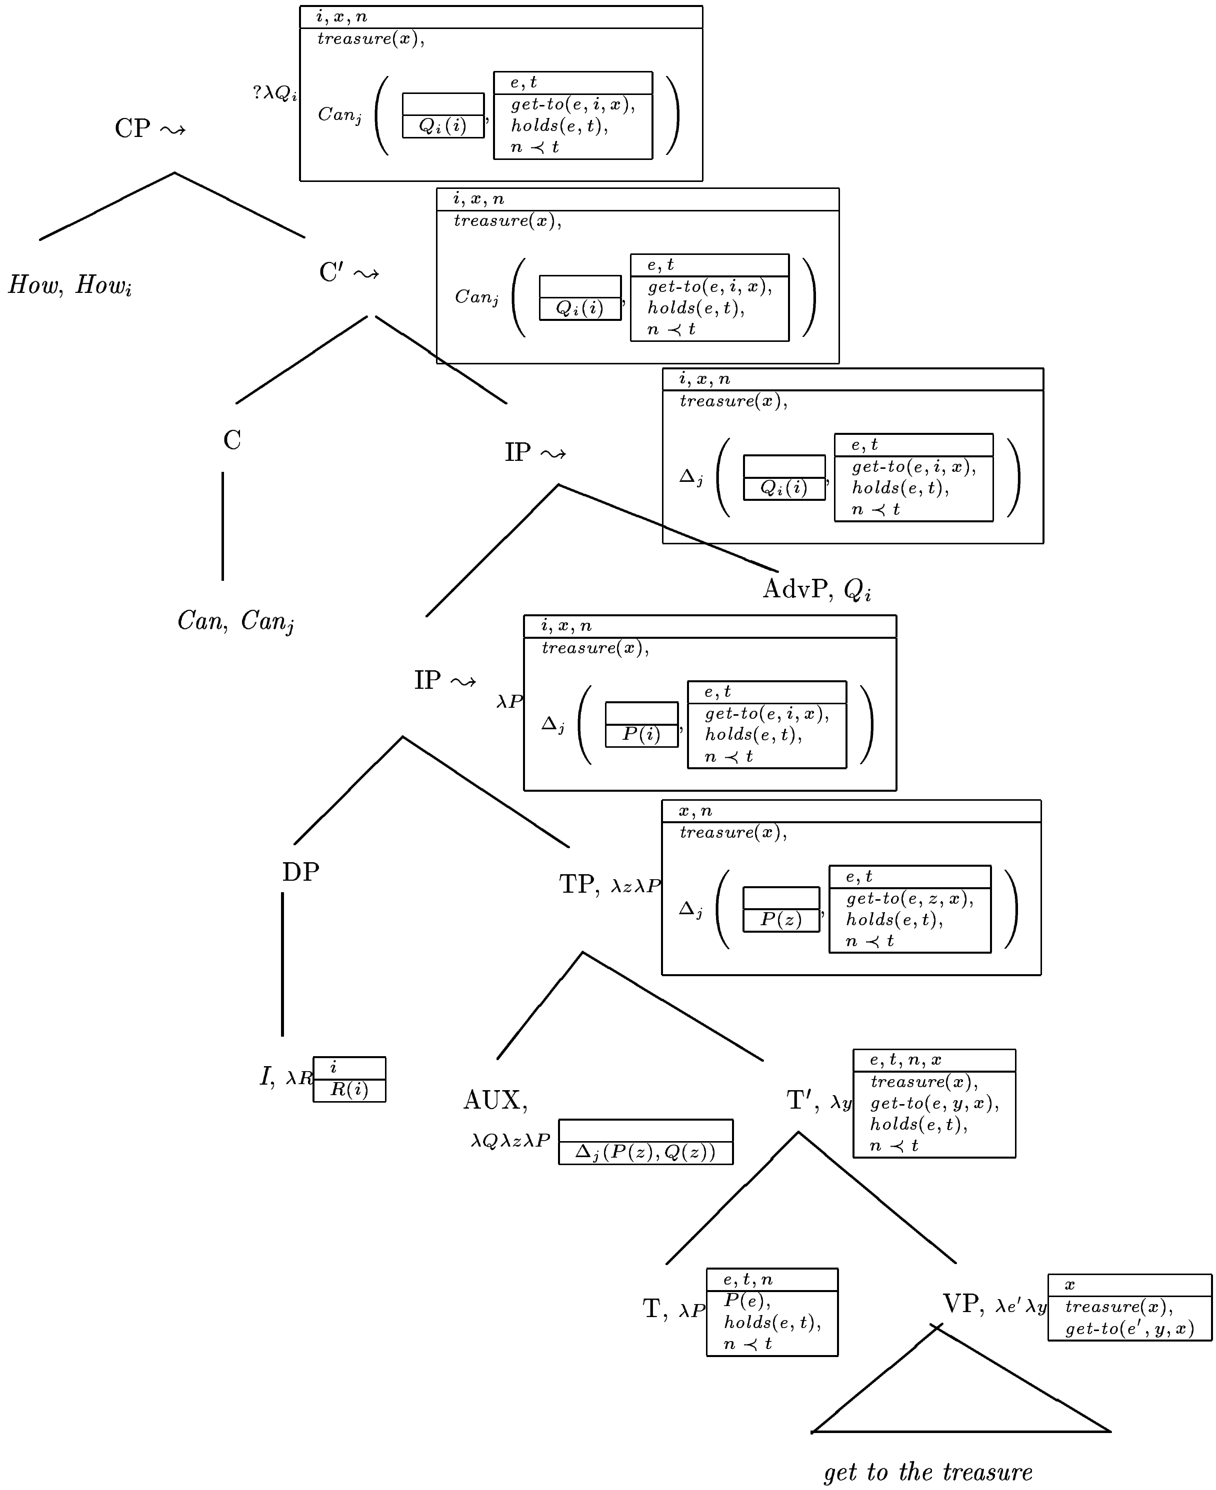

Figure 1: The Bottom up DRS Construction of *How can I get to the treasure?*

6.1 The Question Answer Pair

We assume $QAP(A{:}\alpha, B{:}\beta)$ can hold only if α is a question and β is a (true direct) answer to that question, as determined by the compositional semantics of questions and answers specified above. This in essence is a coherence constraint on QAP; it restricts the semantic relation between the constituents α and β that it connects. For example, if α is an unrestricted *how*-question, then $QAP(A{:}\alpha, B{:}\beta)$ can hold only if β specifies an adverbial of manner, which is a true way of doing the event described in α. In fact, these aren't only necessary conditions, but they are normally sufficient too. More formally, let $Sat\text{-}QAP(\alpha, \beta)$ mean that α and β satisfy the coherence constraints we've just imposed on the semantics of $QAP(\alpha, \beta)$. That is, α is a question and β is a true direct answer to that question.[22] Then if β is to be attached to α, and α and β satisfy the coherence constraints imposed by QAP, then normally α and β are a question answer pair:[23]

- `Question Answer Pair (QAP):`
  $(\langle\tau, \alpha, \beta\rangle \wedge Sat\text{-}QAP(\alpha, \beta)) > QAP(\alpha, \beta)$

We think of $QAP(\alpha, \beta)$ as subordinating β to α, so that follow up questions can be related directly to the previous question α and hence the content of the questions can mutually influence each other. If $QAP(\alpha, \beta)$ were treated as a coordinating relation, then the follow up questions couldn't attach to α since it wouldn't be on the right frontier. It's important to stress that it may be difficult or even impossible to infer $Sat\text{-}QAP(\alpha, \beta)$ with certainty for the agent that asks the question, because he may not know for sure that β is true. However, we note that there may be many methods of nonmonotonically inferring $Sat\text{-}QAP(\alpha, \beta)$, so that the `QAP` axiom fires. For instance, we might suppose that if A asks a *how* question α, and β gives an adverbial of manner that is a doxastically possible direct answer to α—that is, is a direct answer to α in some world compatible with A's beliefs—then that is normally sufficient for B to conclude $Sat\text{-}QAP(\alpha, \beta)$. In other words, we assume the questioner takes the answer he gets to his question to be true, unless he has reason to believe the contrary. This assumption, which we won't formalize here, echoes Grice's maxim of Quality.

With even this minimal account of QAP in SDRT, we can now exploit the compositional semantics and predict the differences between (3) and (4).

(3) a. A: How can I get to 6th Street?
 b. B: You can get there by asking someone Downtown.
 c. A: ?What's his name?

(4) a. A: How can I get to 6th Street?
 b. B: There's someone Downtown that you could ask.
 c. A: What's his name?

By the monotonic procedure for DRS construction, the representation of (3b) and (4b) are respectively (3b′) and (4b′) (before the anaphors are resolved):

[22] It is important to note there may be other coherent responses to a question besides a true direct answer. Indeed, a coherent response may even be another question. We will attach such coherent responses with a rhetorical relation other than QAP, to reflect the fact that they have a different relation to the question than that of a true direct answer. The analysis of coherent responses that aren't true direct answers is examined in detail in section 6.2.

[23] We make this rule a default because there are potential exceptions. For example, consider the situation where A asks *When is the train to Philadelphia?* and B looks at his watch, panics and says *3pm*, meaning the time now is 3pm and he's late. If the rule were not a default, then we may get $QAP(\alpha, \beta)$ in this case (so long as it's true that the train to Philadelphia leaves at 3pm), even though B didn't answer the question. The correct sc sdrs for this situation would be derived from more specific default rules than `QAP` that conflict with it.

(3) b′.

$$\begin{array}{|l|}\hline you, d \\ \hline downtown(d), \\ \\ Can \left(\begin{array}{|l|}\hline e', t', y \\ \hline ask(e', you, y, q), \\ q = ?, \\ location(y, d) \\ \hline \end{array}, \begin{array}{|l|}\hline e, t, x \\ \hline get\text{-}to(e, you, x), \\ holds(e, t), \\ n \prec t, \\ x = ? \\ \hline \end{array} \right) \\ \hline \end{array}$$

(4) b′.

$$\begin{array}{|l|}\hline you, y, d \\ \hline downtown(d), \\ \\ Can \left(\begin{array}{|l|}\hline e', t' \\ \hline ask(e', you, y, q), \\ q = ? \\ \hline \end{array}, \begin{array}{|l|}\hline e, t, x \\ \hline get\text{-}to(e, you, x), \\ holds(e, t), \\ n \prec t, \\ x = ? \\ \hline \end{array} \right) \\ \\ location(y, d) \\ \hline \end{array}$$

In both cases, the conditions $x = ?$ and $q = ?$ stem respectively from the anaphor *there* and the fact that the verb *ask* requires a question that is asked. These anaphors are resolved by reasoning that (3b)/(4b) are related to (3a)/(4a). Indeed, the definition of availability in SDRT means that q and x must pick their antecedents from those that are accessible in DRT terms (i.e., not in embedded DRSs) in the available constituent in the discourse structure (i.e., constituents on the right frontier of that structure). There is only one available constituent: (3a)/(4a). So regardless of the rhetorical relation that connects them, the only candidate for x is *6th Street*, and the only candidate for q is the question *How can I get to 6th Street?*

Different syntactic structures in (3b) and (4b) lead to the discourse referent y (which represents the person to be asked) being available in SDRT terms for anaphoric reference in (4c), but not in (3c). We haven't had to reason with A's or B's cognitive state to predict this. We didn't exploit open ended nonmonotonic inference procedures with cognitive states, such as those proposed in Perrault (1990), to produce the representations of (3b) and (4b). And so we haven't had to consider issues such as: Is B's belief that someone knows how to get to 6th Street *de dicto* or *de re*? Rather, the procedure was monotonic. Thus, modeling the way information is presented in a sentence can refine the constraints on pronominal reference supplied by the more open ended plan recognition techniques. Obviously, DRS structure alone isn't sufficient to define all the relevant constraints on pronominal reference (see Asher (1993) for detailed arguments). But it can play a useful part in the story, because without it one would have to obtain all constraints on anaphora from the intentional structures behind the dialogue. This requires a distinction among plan structures that is too fine grained to be motivated independently.

6.2 Indirect Answers and Partial Answers

Some responses aren't direct answers, but nevertheless contribute in coherent ways as responses to the question.[24] This occurs in (1), where (1b) isn't an adverbial of manner.

(1) a. A: How can I get to the treasure?
    b. B: It's at the secret valley.

[24]Indeed, on a search of two million words of the BNC corpus, 14 questions starting *How can...* were found, and none of these had direct answers in response (i.e., adverbials of manner). However, apart from the two examples (30b) and (30c) cited below, they were all responses like that in (30a), where an adverbial of manner can be inferred.

Such a response won't meet the coherence constraints on *QAP*. However, it is an indirect answer, in the sense that we can compute the direct answer (29) from it, given the discourse context and what we know about the world.

(29) You can get to the treasure by going to the secret valley.

Indirect answers offer another way (but not the only other way) of supplying a coherent response to a question. They are distinct from, but also linked to, direct answers.

Assume for now that (29) is true (in particular, that A can go to the secret valley). Then the DRS representing (29) can attach to the representation α of (1a) with *QAP*. We define an indirect answer to be a response where the interpreter can infer from it a proposition γ such that $QAP(\alpha, \gamma)$ holds. So note that unlike the definition of direct answerhood, indirect answers are interpreter-dependent. Predicting when the interpreter can infer a direct answer γ from the response to an interrogative requires complex reasoning about the discourse context and cognitive state of the interpreter. For example, (1b) is an indirect response to (1a), because A's knowledge about locations, spatial displacements, and actions allows him to compute the direct answer (29) from (1b).

In contrast to (1b), A can't support a belief that B believes that A thinks he can compute a direct answer from the semantic content given in (1b′).

(1) a. A: How do I get to the treasure?

b′. B: ?Mary's hair is black.

To believe this, A would have to assume that B thinks A has sufficient knowledge to compute an adverbial of manner for getting to the treasure from (1b′), and this is untenable. Therefore, A cannot compute a direct answer which B intended to convey by uttering (1b′). Indeed, A believes that B believes that A *cannot* compute a direct answer from the response (1b′). And since indirect answers are constrained to be ones that the questioner can compute direct answers from, A is committed to building his SDRS for B, with some rhetorical relation between (1a) and (1b′) other than the one which stipulates that (1b′) is an indirect answer to the question (1a). But there are no other candidate rhetorical relations by which A can build B's SDRS; A can't work out what B intended by (1b′) at all, given his beliefs about B's KB. And so, by the coherence constraint `Computability` (which stipulates that the interpreter must be able to work out what the speaker intended to convey by his utterance), the response (1b′) is incoherent for A.

We now formalize these intuitions about the connection between direct and indirect answers. We represent the relationship between indirect answers and questions at the discourse structure level via the relation *Indirect Question Answer Pair* (*IQAP*). Now, note that by the axiom `Weak Deduction` and `Believing Inferences` given in section 4.1, if A can nonmonotonically infer γ from his KB with β added to it, then $\mathcal{B}_A(\beta > \gamma)$ holds. So `Axiom on IQAP` captures the coherence constraint we want: if $IQAP(A{:}\alpha, B{:}\beta)$ holds, then there is a (direct) answer γ such that A can nonmonotonically infer it from his KB augmented with the indirect answer β:

(30) a. A: How can we help reduce this kind of accident?
B: For a start, novice pilots should have some experience of aerotowing in single seaters before converting to types which are not fitted with a nose hook, and gliders intended for inexperienced pilots should always be manufactured with a nose hook for aerotowing.

b. A: How can I get rid of head lice?
B: It's very important to get treatment for lice and nits as soon as you find them.

c. A: How can I do that?
B: That is in the hands of the government.

- `Axiom on IQAP`:
  $IQAP(A{:}\alpha, B{:}\beta) \rightarrow \exists\gamma(QAP(A:\alpha,\gamma) \wedge \mathcal{B}_A(\beta > \gamma))$

The axiom on *IQAP* entails that direct answers are also indirect answers; *IQAP* generalizes *QAP* because $\beta > \beta$ is an axiom on DICE and A believes these axioms.[25] More importantly, this axiom links information about A's cognitive state to compositional semantics and the rhetorical relation used in discourse structure. On the one hand, the conditional $\beta > \gamma$ could be believed by A because he inferred γ nonmonotonically from his KB augmented with β, and then used the axiom `Weak Deduction` and `Believing Inferences` to believe $\beta > \gamma$. So A may believe this conditional because of his background knowledge, including his knowledge about plans of action. On the other hand, the rhetorical structure $QAP(A{:}\alpha, \gamma)$ is true only if γ is a direct answer to α, as defined by the compositional semantics of interrogatives and answers we gave in section 5. So our framework is rich enough to express rules which relate three levels of analysis: compositional semantics, rhetorical structure and cognitive states.

This axiom is quite a weak condition on indirect answerhood, because it stipulates only that A be able to nonmonotonically infer a direct answer γ from his *whole* knowledge base, giving him the opportunity to exploit plans and intentions as well as propositions. For example, if the response to a question is an instruction, as in (31), then A can use his knowledge of plans to infer the direct answer *the platform that the porter takes you to*:

(31) a. A: What platform is the 1:20 to Philadelphia leaving from?
     b. B: Follow that porter. He'll put you on the train.

Thus in spite of the fact that (31b) features a directive speech act, which is distinct from the declarative speech act that direct answers to (31a) have, (31b) nevertheless satisfies the above constraint on indirect answerhood. In this sense, speech acts take a back seat in this theory. Note, however, that even though *the platform the porter takes me to* is a direct answer that A can compute from (31b), it does *not* necessarily follow that A *knows* which platform the train is leaving from. This is because A knowing this particular direct answer may not fulfill the minimal demands placed on *knows* by the particular type of knowledge at hand—that is, the interrogative *which platform the train leaves from*—as discussed in section 5.1.

We should stress that this analysis of indirect answerhood is distinct from Groenendijk and Stokhof's in two important respects. First, the connection between an indirect answer and a direct one is *nonmonotonic*. Second, because of the `Weak Deduction` theorem in DICE, the questioner *reflects* on whether he can nonmonotonically compute a direct answer from the indirect one.

Like the constraint on *QAP*, the necessary condition for *IQAP* is normally sufficient too. So if $Sat\text{-}IQAP(A{:}\alpha, B{:}\beta)$ means that α and β satisfy the above coherence constraint on *IQAP* (that is, the consequent of the `Axiom on IQAP` holds for α and β), then the following specifies that in such a situation, one normally attaches the constituents together with *IQAP*:

- `Indirect Question Answer Pair (IQAP)`:
  $(\langle\tau, \alpha, \beta\rangle \wedge Sat\text{-}IQAP(\alpha, \beta)) > IQAP(\alpha, \beta)$

`IQAP` implies that as a byproduct of inferring a rhetorical relation between α and β, the agent computes a direct answer. In (31), for example, A can attach (31b) to (31a) with *IQAP* only if he can compute the direct answer *the platform the porter takes you to* from his KB and the instruction (31b), presumably via his knowledge about causation and action.

[25] Notice that given our axioms, indirect answers and direct answers don't exhaust the set of answer-like responses that we might encounter in dialogue. But we don't intend our classification to be exhaustive.

On the other hand, in a dialogue such as (15), the response B gives doesn't enable the hearer to work out a direct answer, and both A and B know this:

(15) a. A: How do I get Downtown?

b. B: Well, you get there either by taking I-79 or I-279, but I'm not sure which.

As we mentioned in section 5, (15b) is a partial answer to (15a), because it helps the reduce the set of possible propositions that can be direct answers, in the sense of Groenendijk and Stokhof. Under the assumption that A had no prior knowledge at all about how to get Downtown, this particular partial answer (15b) is *not* an indirect answer, because it fails to satisfy the *de re* condition that it entail an adverbial of manner for getting Downtown.[26] In examples like (31) and (15), if A doesn't have the necessary information in his KB to compute a direct answer from B's response, then he can't attach the response with *IQAP*. B's response is coherent for A as we've defined it, but to account for this we will have to suppose that A uses a rhetorical relation other than *QAP* or *IQAP*. The strategy we use for attaching partial answers is similar to that for direct answers and indirect answers. We introduce the rhetorical relation Question Partial Answer Pair (*QPAP*), and place coherence constraints on it which reflect the compositional semantic definition of partial answerhood. That is, $QPAP(\alpha,\beta)$ can be true only if α is a question, and β is a partial answer to that question as we have defined; namely, β reduces the set of possible propositions that can be direct answers to α. Now suppose that $Sat\text{-}QPAP(\alpha,\beta)$ means that α and β satisfy the coherence constraints on *QPAP*. Then `QPAP` below implies that these necessary conditions on *QPAP* are normally sufficient too:

- `Question Partial Answer Pair (QPAP)`
  $(\langle\tau,\alpha,\beta\rangle \wedge Sat\text{-}QPAP(\alpha,\beta)) > QPAP(\alpha,\beta)$

`QPAP` ensures that, by default, *QPAP*(15a, 15b) holds. Furthermore, in the context where A's cognitive state is such that if one adds the compositional semantic information in (15b) to his set of beliefs he can infer a direct answer to his question, then *IQAP*(15a,15b) will be inferred as well via the rule `IQAP`. Otherwise, the *IQAP* relation isn't inferred in this case. There are some contexts where a partial answer provides sufficient information for the questioner that he can achieve the goals that prompted him to ask the question. In other cases, the partial answer won't provide enough information. We discuss what happens in this latter case in section 8.

Let's formally analyse a particular example in detail, using the rules we've introduced so far:

(1) a. A: How can I get to the treasure?

b. B: It's at the secret valley.

Given our semantics of *how*-questions, the DRS representing (1a) is α below. The one representing (1b) is β:

$$(\alpha)\qquad ?\lambda P \begin{array}{|l|}\hline i,n,x \\ \hline treasure(x), \\ Can\left(\begin{array}{|c|}\hline \\ \hline P(i)\\ \hline\end{array}, \begin{array}{|l|}\hline e,t\\ \hline get\text{-}to(e,i,x), \\ holds(e,t),\\ n \prec t\\ \hline\end{array}\right)\\ \hline\end{array}$$

[26] Intuitively, (15b) implicates that B doesn't know a direct answer to (15a). We return to the matter of how B's cognitive state influences the way he responds to questions in section 7.

$$(\beta) \quad \begin{array}{|l|} \hline e', y, z, t', n \\ \hline at(e', z, y), \\ secret\text{-}valley(y), \\ z = \textit{?}, \\ holds(e', t'), \\ n \subseteq t' \\ \hline \end{array}$$

Our task now is to connect them together with a rhetorical relation. The only antecedents accessible to the anaphor z in β are those in α and β. Given the syntactic and semantic agreement constraints on z, the only possible antecedent is x—the treasure. So we can assume that the representation of (1b) is β_1 before we attempt to compute the rhetorical relation.

$$(\beta_1) \quad \begin{array}{|l|} \hline e', x, y, t \\ \hline treasure(x) \\ secret\text{-}valley(y) \\ at(e', x, y) \\ hold(e, now) \\ \hline \end{array}$$

To complete the SDRS for (1), we must compute the rhetorical relation between α and β_1 using DICE. We will consider two cases: one where A knows that he can go to the secret valley, and one where he can't but B thinks he can. Let's examine the first case first. Let's also assume, as intuitions would dictate, that A knows a general, domain level plan which represents the following information: the action of going to the location y of an object x generates the action of getting to the object x (where *generates* is assumed to be defined as in Pollack (1990)). Then using techniques suggested by Pollack (1990) and Lochbaum (1993), one can fill in the x parameter of this plan with the treasure. Moreover, the semantics β_1 of (1b) permit A to fill in the parameter y with the secret valley. So this information obtained from the compositional semantics β_1 of (1b), together with the above very general domain plan, allow A to construct a generates relation between actions: going to the secret valley generates getting to the treasure. Goldman (1970) and Pollack (1990) suggests that the generates relation between actions in plans is closely connected to natural language *by*-phrases. In our terms, this is tantamount to the assumption that if $\alpha_1(x)$ generates $\alpha_2(x)$ in a domain plan and x can do α_1, then $Can(\alpha_1(x), \alpha_2(x))$ holds—that is, x can do α_2 by doing α_1. So, assuming this link between the modality *Can* in our extension of DRT and the definition of generation in plans which Pollack alludes to, A uses this generates relation between actions that he has inferred via the semantic content of β_1, to infer nonmonotonically the direct answer (29) to (1a), as represented in the DRS (29′) given earlier. Let's rename (29′) as γ. Then this discussion about A's KB demonstrates that A can use the `Weak Deduction` axiom in DICE to infer a new condition: $\beta_1 > \gamma$.

Since we are for the moment assuming that A can go to the secret valley and that A believes this (thus assuming γ is true), A can infer $QAP(\alpha, \gamma)$ via Defeasible Modus Ponens (DMP) on `QAP`. So $\beta_1 > \gamma$ and $QAP(\alpha, \gamma)$ are now part of his KB. Hence the coherence constraints on $IQAP(\alpha, \beta_1)$ are satisfied in his KB. Thus the antecedent to `IQAP` is verified in his KB. The antecedent to `Background` is also verified, since β_1 describes a state. The consequents of both these `IQAP` and `Background` are consistent with his KB and with each other. Therefore, by defeasible modus ponens (DMP), A nonmonotonically infers $IQAP(\alpha, \beta_1)$ and $Background(\alpha, \beta_1)$.

It is important to stress that inferring this discourse structure in this example crucially depended on the assumptions both that A can go to the secret valley and that A believes this. If A can't go to the secret valley or doesn't believe it, then by clause (b) of the semantics of *Can*, A thinks γ is false, and so can't constitute part of an answer to the question (since $QAP(\alpha, \gamma)$ entails γ is true). This restriction matches intuitions. B's attempts to perform a speech act of providing a correct answer to α fail if A thinks the associated direct answer is false. So if A doesn't believe that he can go to the secret valley, then he cannot accept an SDRS that stipulates that the speech act type *IQAP* has been performed.

This raises the question: what discourse structure does A compute in this second case to be considered? And on what grounds is (1b) a coherent response, given that A can't compute that it's an answer, indirect or otherwise? We will return to this question in section 8, and demonstrate how we capture the fact that (1b) is a coherent, if unsatisfactory, response to A's question. But before we do this, we need to examine the link between questions and responses in discourse, and the way an interpreter's interpretation of the discourse leads him to certain beliefs about what the participants believe and intend.

7 Reasoning with Cognitive States and Discourse Structure

The discourse relations *QAP* and *IQAP* that we've discussed in the previous two sections affect what the interpreter thinks about the speaker in systematic ways. To get at these effects, we will use the principles for reasoning about about cognitive states that we introduced in section 4.2. We will ultimately exploit these effects to model how the participants in the dialogue negotiate disagreements in assumptions made about what they believe and know. And this will be important when computing the SDRS that A builds for (1), when A believes she can't get to the secret valley.

The reason question asking and answering have systematic cognitive effects is that like most speech acts, question asking and question answering are intentional. By looking at the cognitive assumptions behind question asking, we will be able to derive the cognitive effects of answering using our principles of cooperativity and the Practical Syllogism.

Typically, the questioner has a goal in mind in asking the question. People ask questions for many purposes, but presumably one purpose that always *prima facie* applies in cooperative dialogue—which is the only kind of dialogue we consider here—is the purpose to get a correct (i.e., true) answer to the question. A cooperative respondent should typically respond with a correct (direct or indirect) answer as far as he is able to judge. So as Searle (1969) suggests, questions are a kind of request. But behind requests and questions, there lies a certain desire or goal—for questions, this is a *question related goal* or QRG

Why would A ask the question? We cannot say unequivocally; there might be all sorts of reasons. But as we've stated, one reason for A asking $?\alpha$ is usually that A wants to know an answer, and doesn't have this knowledge now. We can formalize this in our logical framework as `QRG`:

- `QRG`: $A{:}?\alpha > \exists\gamma(QAP(\alpha,\gamma) \wedge \mathcal{W}_A\mathcal{K}_A(\gamma) \wedge \neg\mathcal{K}_A(\gamma))$

This rule bears a close resemblance to Searle's (1969) preparatory condition on queries. However, his preparatory conditions are too strong, and we have weakened them by making this rule a default. For example, his preparatory conditions are problematic in cases such as exams. But being in an exam situation is a special, more specific situation, which will override the above default `QRG`.

If we use `QRG` to reason (defeasibly) about the QRG of questions, then we see that A's asking, for example, a question α of the form *How can I ϕ?* entails defeasibly that there's a direct answer χ to α, such that $\mathcal{W}_A(\mathcal{K}_A(\chi))$ and $\neg\mathcal{K}_A(\chi))$ are true. Clause (a) of the `Practical Syllogism` is thus verified with $\mathcal{K}_A(\chi)$ replacing ψ (and A replacing B), and if A believes that by asking a question α he can get an answer then clause (b) is verified too, with $A{:}?\alpha$ replacing ϕ. So we derive by `Practical Syllogism` that $\mathcal{I}_A(A{:}?\alpha)$. That is, we have derived within DICE A's intention to perform the speech act of asking a question.

We can go further and tie question answering behavior to further goals. Suppose, for instance, that if an agent asks a *how*-question of the form *How can I ϕ?*, normally he wants to achieve ϕ but believes that $\neg\phi$. Suppose further that B answers A's question with β, and that β suffices for A to compute a true direct answer to his question. That is, β tells A that there is a plan ψ that A can do that will normally bring about ϕ. Once again the antecedent of the `Practical Syllogism` is verified. So, we can conclude $\mathcal{I}_A(\psi)$; that is, A intends to execute the plan ψ given by the answer. Thus the `Practical Syllogism` is used here to reason about how A's interpretation of the dialogue affects the actions he intends to do.

A concrete example of this line of reasoning surfaces in (1). We'll assume, in line with the above, that in addition to the QRG given by `QRG` (i.e., the goal of knowing a correct answer to (1a)), A also has the QRG given by the DRS α', which represents *A gets to the treasure*. Furthermore, let's suppose for now that A computes the SDRS $IQAP(\alpha, \beta_1)$ for the DRS representations α and β_1 of (1a) and (1b) that we defined earlier. Hence A can infer from β_1 that for some ψ, $Can(\psi, \alpha')$ holds and therefore by the semantics of *Can*, he can do ψ, and doing it nonmonotonically will yield α', as long as he intends α'. So assuming nothing else about A, we infer in virtue of the `Practical Syllogism` that A intends ψ. Thus, we have the following fact: let P_a be the plan that specifies the manner a for the question q which is of the form *How can I ϕ?* Then: $(How\text{-}can\text{-}I\text{-}\phi(q) \wedge QAP(A{:}q, a)) \approx\!\!\!\mid \mathcal{I}_A(P_a)$ is valid. Note furthermore, that B's response (1b) was *appropriate* in this case, since it enabled A to achieve his QRG of getting to the treasure. It did this, because it provided enough semantic content for A to compute an adverbial of manner that he can do for getting to the treasure. That is, it was appropriate in this case, precisely because A could compute the rhetorical relation $IQAP$.

The particular specification of the intended plan will depend on what plans were already in A's KB, as we stated in section 6. These considerations are discussed in detail by Litman and Allen (1990) among others, and are beyond the scope of this paper. We would need to say a great deal more about plan generation and plan revision. But we still have shown how the `Practical Syllogism` could be used to affect intentions. And in this sense, we refine Ginzburg's (1995) analysis of questions and answerhood, in that we show how questions and answers interact with goals and intentions.

Let's now use the `Practical Syllogism` to look at how responses work. If B's cooperating, then he wants to do more than answer the question correctly. He also wants to help A achieve whichever other goals lay behind the question. In other words, if A has a QRG p—and so in particular, $\mathcal{W}_A p$ and $\mathcal{B}_A \neg p$ hold—then normally B cooperates by adopting p as his goal too. This is captured by the axiom `Cooperation` that we gave in section 4.2.

This can be used to infer that B thought he was answering the question appropriately. That is, B thought his answer had sufficient semantic content, that A would be able to achieve his QRGs.[27] For suppose that B reasons that A's QRG for α is p. Then $\mathcal{W}_A p$ and $\mathcal{B}_A \neg p$ form part of B's KB. So by DMP on `Cooperation`, B infers $\mathcal{W}_B p$ and $\mathcal{B}_B \neg p$. Let ϕ in the `Practical Syllogism` stand for the discourse action $\langle A{:}\alpha, A{:}\alpha, B{:}\beta\rangle \wedge Info(\alpha, \beta)$, where $Info(\alpha, \beta)$ stands for the translation μ into the glue logic of the compositional semantic content of α and β, and let ψ stand for p. Then because of the above inference from `Cooperation`, clause (a) in the antecedent of the `Practical Syllogism` holds. Moreover we assume that B's discourse action was intentional, and so the consequent (c) also holds. So as in Asher and Lascarides (1994), the antecedent (b) can be abduced: B believes that his discourse action nonmonotonically achieves p. That is, A will fulfil his goal p as a result of B's discourse action, by default. Since by `QRG` p by default is to learn a direct answer, B believes that his discourse action nonmonotonically achieves: A can infer a direct answer γ.

The reasoning that we have just put in the head of B is reasoning that A can also do *about B*, and presumably would do about B assuming `Cooperation` and the fact that B responds to A's question with β. So (32) holds in our theory:

[27] In Ginzburg's terms, B thought he had fully resolved A's question for A. This fills a gap in Ginzburg's analysis, since we *derive* this information about what B normally thinks his dialogue move is in responding.

(32) $\texttt{Practical Syllogism}, \texttt{Cooperation}, \texttt{QRG}, \langle\tau, A\!: ?\alpha, B\!:\!\beta\rangle \mathrel{|\!\approx} \mathcal{B}_B(\mathit{IQAP}(A\!:\!?\alpha, B\!:\!\beta))$

That is, the general axioms linking belief, desire and intention, and the axiom that normally, one QRG for asking a question is to get a (direct or indirect) answer, allow one to infer nonmonotonically that B's response β to a A's question α is one that B believed would be an indirect answer.

Now, since `Practical Syllogism`, `Cooperation` are axioms of DICE and since we cannot derive nonmonotonically $\mathcal{B}_B(\mathit{IQAP}(A\!:\!?\alpha, B\!:\!\beta))$ if we do not assume $\langle\tau, A\!: ?\alpha, B\!:\!\beta\rangle$, we have by the `Weak Deduction` rule of CE given in section 4.1 the following fact:

- DICE $\mathrel{|\!\approx} \langle\tau, A\!:\!?\alpha, B\!:\!\beta\rangle > \mathcal{B}_B(\mathit{IQAP}(?\alpha, \beta))$

In other words, the default rule `Answer` is an axiom of DICE:[28]

- `Answer`: $\langle\tau, A\!:\!?\alpha, B\!:\!\beta\rangle > \mathcal{B}_B(\mathit{IQAP}(?\alpha, \beta))$

This cooperative principle that normally, people think their responses are indirect answers to the question, has been derived from general axioms about how beliefs and desires affect our intentions, and about the goals that normally lie behind an question. These more general axioms have utility in all kinds of situations, not just the ones where dialogue is taking place.

The default `Answer` may seem like a strong axiom. However, it is a default, and so any situation where there is more specific conflicting information about the rhetorical role of the response will override it. Indeed, such situations may be relatively common.[29] For example, the original interrogative may generate the need for clarification, and in such a situation, a conflicting default with a more specific antecedent than `Answer`'s will override it. For example, consider the following dialogue:[30]

(33) a. A: Where is the faucet?
b. B: Most of us get our drinks from the fridge, because the water isn't safe here.
c. B: ?There is a woollen shop on the High Street.

Let's assume that the compositional semantics of (33b) provides sufficient information that both A and B can conclude nonmonotonically that B believes that there is no direct answer γ to A's question that will help A fulfill his QRG. Then in this dialogue situation, the fact (32) holds, but so does the following (more specific) fact (34):

(34) $\texttt{Practical Syllogism}, \texttt{Cooperation}, \texttt{QRG}, \langle\tau, A\!:\!?\alpha, B\!:\!\beta\rangle, \mathit{Info}(?\alpha, \beta) \mathrel{|\!\not\approx} \mathcal{B}_B(\mathit{IQAP}(\alpha, \beta))$

We *don't* get a nonmonotonic entailment that $\mathcal{B}_B \mathit{IQAP}(\alpha, \beta)$ when the premises are those given in (34), because the `Practical Syllogism` won't apply with ϕ in the `Practical Syllogism` being replaced with a direct answer γ. Therefore, A and B can both reason that `Answer` is verified, but doesn't apply in the context of the dialogue (33) because of the more specific conflicting information. We assume that in cases where B believes there's no direct answer that fulfills the questioner's QRGs, some relation other than *IQAP* is used to attach the response to the question. In contrast, one cannot attach utterance (33b′) to the question with a rhetorical relation, because

[28] Note that this is just a default, and there are many circumstances under which the antecedent of this law may hold, and the consequent not. For example, if B doesn't know an answer to the question, or if he doesn't have time to answer etc.

[29] It is important to stress a distinction between reasoning by default, and reasoning by frequency. The default axioms of DICE are axioms for reasoning when you don't have information to the contrary; they do *not* necessarily encode what is most frequently the case.

[30] Thanks to an anonymous reviewer for suggesting this dialogue.

in contrast to (33b), (33b′) doesn't contain enough information that one can conclude that B thinks there are no direct answers γ that fulfill A's QRGs, and yet A can't infer it's an indirect answer either. Indeed, A cannot come to any nonmonotonic conclusions about what B intended by his utterance, and so `Computability` is violated, leading to incoherence.

We have shown that the axioms we've given validate defaults like `Answer`, which enables the questioner A to come to nonmonotonic conclusions about what B believes, on the basis of what he's said. One can go further in reasoning about cognitive states from what's been said in the discourse. One can use `Answer` to infer what A thinks B thinks A thinks. For suppose A uses `Answer` to infer his model of B's SDRS. Then by the `Axiom on IQAP`, A believes that B believes that his response β was one which A could use to derive a direct answer γ. So A learns something about what B thinks A thinks. A thinks that $\exists\gamma(QAP(\alpha,\gamma) \wedge \mathcal{B}_A(\beta > \gamma))$ is in B's KB. So if it wasn't already there, A adds this to his model of B's KB, as a direct consequence of using `Answer` to compute what B's SDRS is. A could even take things further, and reason about exactly what direct answer B thinks A will compute, thus turning the above *de dicto* belief $\mathcal{B}_B(\exists\gamma(QAP(\alpha,\gamma) \wedge \mathcal{B}_A(\beta > \gamma))$, which is in A's KB, into the *de re* belief $\exists\gamma(QAP(\alpha,\gamma) \wedge \mathcal{B}_B\mathcal{B}_A(\beta > \gamma))$. In the case of (1), for example, A can compute the *de re* belief in his model of B's KB, by assuming B has modeled A's knowledge of geometry.

Finally, our axioms about the interaction between cognitive states and discourse structure validate the intuition that when B thinks he's provided an answer to a question, he thinks that the questioner A thinks so too. Of course, the respondent B may use axioms and theorems other than `Answer` to infer $IQAP(\alpha,\beta)$ in his model of the dialogue—he may use `IQAP`, for example. But at any rate, whichever way B infers $IQAP(\alpha,\beta)$, the constraints on $IQAP$ given in `Axiom on IQAP` mean that whenever B infers it, $\exists\gamma(QAP(\alpha,\gamma) \wedge \mathcal{B}_A(\beta > \gamma))$ must be in B's KB. Moreover, this belief affects the way B builds his model of A's SDRS. DICE is mutually believed, and so B can compute from his model of A's KB, which contains $\beta > \gamma$ where $QAP(\alpha,\gamma)$, that A uses the default DICE axiom `IQAP` to nonmonotonically infer $IQAP(\alpha,\beta)$. Hence B nonmonotonically concludes $\mathcal{B}_A(IQAP(A{:}\alpha, B{:}\beta))$. This nonmonotonic inference followed directly from $IQAP(A{:}\alpha, B{:}\beta)$ being in B's model of his own SDRS. So, by `Believing Inferences`, $\mathcal{B}_B(IQAP(A{:}\alpha, B{:}\beta))$ entails $\mathcal{B}_B\mathcal{B}_A(IQAP(A{:}\alpha, B{:}\beta))$ in DICE. That is, whenever B believes he's giving an indirect answer to A's question, he believes that A believes this too.

In this section, we have seen how our discourse relations between questions and responses in dialogue invite particular conclusions about the cognitive states of the participants in the dialogue. Our axioms that specify how to compute rhetorical relations , and very general axioms concerning cognitive states and the principles of cooperativity, model how hearers can use what's said in the dialogue as clues to infer things about what the speakers believe about themselves and the hearers. Like Ginzburg (1995), we believe that answering questions appropriately is inherently linked to the speaker's goals and beliefs. This notion is also familiar from plan recognition work such as Lochbaum's (1995). But we have gone further than Ginzburg in the story, in that we have provided some axioms that help us track how beliefs and goals change as the dialogue progresses. And we have provided insights into the plan recognition work, in that we have not had to depend on a very fine-grained classification of plan operators, nor on theory dependent axiomatisations of discourse segmentation, in order to relate the participants' plans and goals to what they say in the dialogue.

8 When The Response Doesn't Provide Enough Information

As we've mentioned in section 4.2, (1b) is coherent for A only if he can compute how β_1 attaches to α in his SDRS, and how it attaches in his model of B's SDRS—as demanded by the attachment constraint `Computability`. So far, we've seen what happens in A's SDRS when A can get to the secret valley and believes he can. We've also seen how this situation leads to certain beliefs in A about B and vice-versa. But now we have to ask, what happens when A believes he can't get to the secret valley?

Let's deal with A's model of B's SDRS first. In the previous section, we derived `Answer`, which describes cooperative behaviour during communication, as an axiom of DICE from principles about how cognitive states interact with discourse structure. As we stated, there are exceptions to `Answer` such as (15) and (35), but they're usually linguistically explicit.

(35) A: How do I get to the treasure?
B: I'm afraid I can't answer that question.

Examples like (15) and (35) would be accounted for by conflicting, more specific default rules than `Answer`, which would enable the participants to infer via Specificity a different rhetorical relation from *IQAP*. So back to (1). `Answer` permits A to satisfy the `Computability` constraint when attaching (1b) to (1a): He uses it to infer that B's SDRS is $IQAP(\alpha, \beta_1)$.

But what about A's own SDRS? In the case under consideration, A believes he can't go to the secret valley. So given our semantics for *can*, he can't use his KB to derive a manner adverbial and so he can't attach β_1 to α with *IQAP* in his SDRS, even though he can infer that this is what B does. Nevertheless, β_1 is at least a partial answer, since it reduces the space of possible answers for A.

So far so good. But there is something more than simple partial answerhood at stake here. In relation to what A would like and what B thinks he has done, the partial answer isn't sufficient. Intuitively, β_1 doesn't provide enough information for A to construct a true answer to the question α, which is what he wants and what B is trying to help him achieve. So A should use the relation *NEI* (*Not Enough Information*) to relate β_1 to α, as well as *QPAP*.

From the standpoint of monologue *NEI* might appear to be an odd discourse relation, because it seems to assert paradoxically the lack of a coherence relation between two constituents. But this would be a misunderstanding of the relation. *NEI* is an acceptable discourse relation for dialogue, because it reflects and codifies in the SDRS a coherent and common move in dialogue—one person's answer to another's question is signalled as not being sufficient.

NEI is inferrable in DICE with the help of the following axiom `NEI`. It states: if B's response β is to be attached to A's question α, and $IQAP(\alpha, \beta)$ doesn't hold although B believes it does, then normally β is to be attached to α with *NEI*. That is, B's response doesn't provide enough information for the questioner A:

- `NEI`:
  $(\langle \tau, A{:}?\alpha, B{:}\beta \rangle \wedge \neg IQAP(\alpha, \beta) \wedge \mathcal{B}_B(IQAP(\alpha, \beta))) > NEI(\alpha, \beta)$

What this axiom makes plain is that *NEI* is invoked when there is a mismatch between what participants in the dialogue

So A has constructed different SDRSs for himself and for B. He uses *IQAP* via `Answer` for B, but *NEI* for himself because he believes γ (or (29′)) is false. A similar analysis would apply to (31), if B believes that the porter will take him to the wrong platform, for example. We'll show in section 9 that the DICE rules we've given capture the intuition that the difference in the SDRSs stemmed from B making incorrect assumptions about A's cognitive state, and we'll see how A then uses (1c) to get B to correct those assumptions.

But first let's contrast the above analysis of (1a,b) with that of (1a,b′).

(1) a. How do I get to the treasure?

b′. ?Mary's hair is black.

Let's assume, as intuitions would dictate, that in contrast to (1b), A is unable to assume that B has a model of A's KB which includes a domain plan that permits A to compute an adverbial of manner from (1b′) for getting to the treasure. For otherwise, A has a very unusual model of what B thinks A can do. Then, A must assume that B believes the consequent of the `Axiom on IQAP` is false. Therefore, by the logic of belief, A assumes B doesn't attach (1b′) to (1a) with *IQAP*; for otherwise A is committed to B having an inconsistent KB. Therefore, although the antecedent to `Answer` is verified, its consequent can't be inferred, because it's inconsistent with A's KB—specifically his model of B's KB. Indeed, A can't come to any conclusions about which discourse relation B intended A to use to attach (1b′) to (1a). For just as in (10), although A can assume that B thinks A believes the antecedent to `Background` is also verified, A must also assume that B believes A believes that its consequent isn't inferrable because the constraint that there is a distinct common topic would be violated. Therefore, the SDRT attachment constraint `Computability` is violated. A cannot infer, even nonmonotonically, which speech act B intended A should recognise. And since `Computability` is violated, the discourse is incoherent for A (although, of course, it may be coherent for B if he does indeed have unusual assumptions about A cognitive state that A doesn't know about).

Another example of a failure to provide even an indirect answer to a question occurs in (2). Just as in dialogue (1), B infers $IQAP(2a, 2b)$ and A infers $NEI(2a, 2b)$ in (2):

(2) a. A: How do I install the assembly?

b. B: Typically, you put the assembly on before tightening the screws on the assembly.

c. A: But I tried that and it was too difficult.

d. B: So let's tighten the screws on the assembly first then.

The analysis proceeds along similar lines to those above. B's response isn't an adverbial of manner, and so it doesn't satisfy the constraints for attaching (2b) to (2a) with *QAP*. However, it refers to a plan which specifies an order between action types. On the assumption that A successfully infers that the plan referred to is part of a plan for installing the assembly, A can infer an adverbial of manner from the response B gives in (2b): Install the assembly by putting the assembly on before tightening the screws on the assembly. But as in (1), this direct answer is false, because he knows that he can't do the adverbial of manner suggested and so the axioms on the semantics of the modality *do* are violated. Nevertheless, as in (1), he uses DMP on `Answer` to infer that B has $IQAP(2a, 2b)$ in his SDRS, and he uses DMP on `NEI` to infer $NEI(2a, 2b)$ for himself. Having detected the discrepancy in their SDRSs, A must tailor his next utterance to informing B that there is this discrepancy.

Our approach also helps relieve the pressure on domain knowledge for constructing discourse structures for dialogue. To see this, let's reconsider (5).

(5)	a.	A: I need to catch the 1:20 to Philadelphia.	α
	b.	Where's it leaving from?	β
	c.	B: Platform 7.	γ
	d.	A: Where do I get a ticket?	δ
	e.	B: From the booth at the far right end of the hall.	ϵ

The compositional and lexical semantics of (5a,b) inform us that A states an intention and then requests information. Since (5a,b) each give rise to a constituent in SDRT—call them α and β

respectively—we need to compute a rhetorical relation between them. Intuitively, the agent is signalling with his question that he does not have enough information to carry out the intention in α. We mark this again with the discourse relation *Not Enough Information* or *NEI*. The DICE axiom `NEI II` below thus reflects this intuition, where $[\mathcal{I}_A(\phi)](\alpha)$ stands for the condition that the SDRS α contain the condition that A intends to ϕ. This must therefore be verified by compositional and lexical semantic information:

- `NEI II`:
  $(\langle \tau, A{:}\alpha, A{:}?\beta \rangle \wedge [\mathcal{I}_A(\phi)](\alpha)) > NEI(\alpha, \beta)$

We understand *NEI*(α, β) as subordinating β to α, to reflect the fact that a coherent response to β must help A produce some action relevant to the realization of the intention stated in α. There may be several occurrences of the *NEI* relation in dialogue, but eventually we will 'pop' out from these, once the questioner has enough information from all his queries to bring about his desire stated in α.

In dialogue (5), both A and B infer *NEI*(α, β) via DMP on `NEI II`. And so the discourse structure at this stage looks as follows:

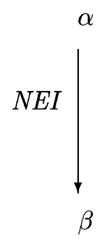

This contrasts with the plan recognition analysis of (5a,b), where commonsense plans are used to compute the relation in the discourse structure. We have used compositional semantics as the clue for computing *NEI*, rather than a commonsense plan, since the antecedent of `NEI II` is verified by compositional and lexical semantic information.

Now, we can compute via the compositional semantics of questions and answers that the constituent γ derived from B's response (5c) is a direct answer to the question β. Moreover, the relation *NEI*$(A{:}\alpha, A{:}?\beta)$ can serve to constrain coherent responses γ to the question β. Intuitively, γ is coherent (for A) only if it helps him achieve the intention in α that prompted the question β. In other words, γ must elaborate a plan which when executed achieves the intention—that A catch the 1:20 train to Philadelphia.

Let's specify this coherence constraint on γ more precisely. Let *Plan-Elab* be the rhetorical relation Plan Elaboration, so *Plan-Elab*(ϕ, γ) means that γ describes part of a plan which, when executed, results in ϕ. Then `Constraint on Answers` encapsulates the relevant constraint. It states the following. If β is a question that A asked to gain information on how to achieve his intention ϕ which is linguistically explicit in α (as asserted by $[\mathcal{I}_A(\phi)](\alpha) \wedge NEI(\alpha, \beta)$); and γ is a direct answer to β (as asserted by $QAP(A{:}\beta, B{:}\gamma)$); then A must be able to use the background knowledge in his KB (which includes the plan library), augmented with ϕ and γ, to nonmonotonically infer that the answer γ elaborates a plan for achieving ϕ.

- `Constraint on Answers`:
  $([\mathcal{I}_A(\phi)](\alpha) \wedge NEI(\alpha, \beta) \wedge QAP(A{:}\beta, B{:}\gamma)) \rightarrow \mathcal{B}_A((\phi \wedge \gamma) > \textit{Plan-Elab}(\phi, \gamma))$

This rule resembles axioms from the AI planning literature on dialogue according to which answers must help the questioner fill in the plans that lay behind his question. However, there are important differences. `Constraint on Answers` is an axiom about rhetorical relations, and the antecedent will be verified on the basis of *linguistic* clues (such as the use of the word *need* in A's utterance), and not clues coming from reasoning about plan recognition.

Let's see how `Constraint on Answers` affects the analysis of (5). The direct answer γ must satisfy `Constraint on Answers`, given that the question β stands in the relation *NEI* to α. In other words, A must believe the conditional $(\phi \wedge \gamma) > \mathit{Plan\text{-}Elab}(\phi, \gamma)$. He can believe this as a result of using the `Weak Deduction` axiom, so long as his KB (which includes background knowledge) augmented with the information ϕ and γ allows him to nonmonotonically infer that γ elaborates a plan for achieving ϕ. So, to check the coherence of γ, A needs to appeal to domain knowledge about plans. But this should be no surprise: it is the same strategy as already found in DICE for inferring *Elaborations* (Lascarides and Asher, 1993). Intuitively, given A's domain knowledge about how to catch a train (specifically, the plan we specified in section 3), the information in γ is part of what is needed to catch a train. More formally, the information in γ is a discourse-permissible subtype ($\mathit{Subtype}_D$) of the information needed to catch a train because of this commonsense plan. So by the default rule given below (cf. the axiom `Elaboration` in Asher and Lascarides 1995), A is able to nonmonotonically infer $\mathit{Plan\text{-}Elab}(\beta, \gamma)$ from his KB when processing dialogue (5), as required by `Constraint on Answers`.

- `Plan Elaboration:`
  $(\langle \tau, \alpha, \beta \rangle \wedge \mathit{Subtype}_D(\beta, \alpha)) > \mathit{Plan\text{-}Elab}(\alpha, \beta)$

Thus, γ successfully attaches to the available constituents on the right frontier: the relations are $\mathit{QAP}(\beta, \gamma)$ and $\mathit{Plan\text{-}Elab}(\phi, \gamma)$, where ϕ is $\mathit{catch\text{-}train}(A)$ and is part of the information in α:

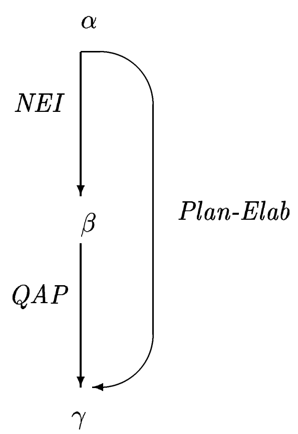

QAP is a subordinating relation, and so the follow up question (5d) can attach to γ, β or α— they're all on the right frontier. The representation δ of (5d) attaches to α in the same way as β did because of `NEI II`. Moreover, we can encode a rule in DICE which models the intuition that when two questions are both asked to achieve a particular intention in the discourse structure, then these questions are related by *Continuation*. Finally, B's response ϵ yields similar inferences to those used to attach γ. So the SDRS that A and B build for dialogue (5) is as follows:

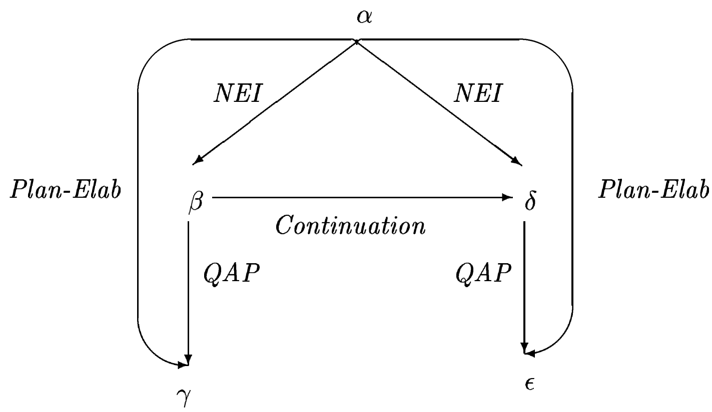

This analysis shows that in DICE, domain knowledge and knowledge about plans are still useful when inferring elaborations. With elaborations in monologue, we may have to use domain knowledge to conclude that one event is possibly part of another event; here, we use knowledge about plans to check that the information that A receives is part of what is needed to carry out a plan to catch a train. However, semantics lets us use this knowledge selectively. We don't always need to use commonsense plans as clues for structuring discourse. The structure of the commonsense plan was used to infer the *Plan-Elab* relations in the above, but not the *NEI*, *QAP* or *Continuation* relations, which also contribute to the overall discourse structure. These latter relations were inferred using only compositional semantics and previous discourse structure as clues.

These examples demonstrate the advantages of using compositional semantics to help provide precise coherence constraints to responses to questions in dialogue. It is unclear how one defines incoherence in the analyses of dialogue which use plan recognition techniques alone (e.g., Lochbaum 1993). Clearly, (1b′) fails to fill in any of the parameters in plans that are in the library and relevant to getting to the treasure. However, this alone cannot be grounds for incoherence, because if it were then it would not be possible to learn new plans from dialogue. Rather, the interpreter at this stage should construct a partial plan which has getting to the treasure as its goal, and which has a partially specified action as a part which utilizes the color of Mary's hair. Of course, this latter action would be almost completely unspecified, because A can't work out how to utilize Mary's hair color to achieve getting to the treasure. So the question then arises: at what stage is such a plan so partial that the dialogue sounds odd? It is not clear what the answer should be if all one has recourse to is the structure of plans and the cognitive state of the participants. However, using the compositional semantics of interrogatives and answers supplies a precise definition of this threshold: it is 'too partial' if it fails to allow A to compute what direct answer to the question B intended to allude to. In the case of *how*-questions, therefore, the plan constructed is "too partial" if it fails to support an adverbial of manner by which to bring about the event in the question. Thus the compositional semantics of interrogatives and answers complements the analyses of dialogues that exploit reasoning about plans, by providing precise specifications of when the plan constructed through interpreting dialogues yields incoherence.[31]

9 Interactions between the Discourse Relations

In the previous sections, we saw how the discourse relations *QAP* and *IQAP* serve to specify constraints on the coherence of responses to questions, and allow one to compute information about the cognitive state of the participants in the dialogue. Here we consider how these relations

[31] Since the compositional semantics of *how*-interrogatives are defined in terms of adverbials of manner, and not all languages contain the same adverbials of manner, nor the same relations to, for example, the *generates* relation in plans discussed by Pollack (1990), the circumstances under which a discourse is coherent are predicted to be language specific. This is a matter for future investigation.

and their constraints engender interactions with other discourse relations like *Contrast*. We will show how the dialogue setting enables us to use *Contrast* to learn about the participant's cognitive states as well.

In section 7, we showed how cooperative participants in dialogue can use what's said to infer things about what they believe about themselves and each other. In particular, we derived axioms such as `Answer` from general principles about the way beliefs, desires and intentions interact in cooperative dialogues: normally, respondents think they're answering the question.

Our DICE axioms also enable an interpreter to detect discrepancies between his beliefs and the speaker's. This is essential, if one is to tailor subsequent utterances to the needs of the participants (Moore and Paris, 1993). This interaction between the SDRSs and cognitive model will lead us to rethink how rhetorical relations may operate in the dialogue. We will focus on *Contrast* and the role of *but* in our examples (1) and (2).

Since DICE enables the interpreter to infer things about the speaker's beliefs, he can utter things to rectify discrepancies that he detects. This is what happens in dialogues (1) and (2). Let's turn back to the analysis of (1a,b), in the situation where A can't go to the secret valley, and examine in more detail how A infers things about B's cognitive state from his utterance (1b).

(1) a. A: How can I get to the treasure?
    b. B: It's at the secret valley.

As we showed in section 8, A computes $NEI(\alpha, \beta_1)$ in his SDRS, and $IQAP(\alpha, \beta_1)$ in his model of B's SDRS. Thus the logical closure of belief and `Axiom on IQAP` mean that the DRS K_B below is part of A's model of B's cognitive model—that is, $\mathcal{B}_A\mathcal{B}_B(K_B)$ holds in A's KB—where α' is the DRS representing *I get to the treasure*, and $P(A)$ is the adverbial of manner that A thinks B thinks A can compute from his answer (1b).

$$(K_B) \quad \begin{array}{|c|}\hline \\ \hline \mathcal{B}_A \left(\begin{array}{c} \beta_1, \\ IQAP(\alpha, \beta_1), \\ Can\left(\begin{array}{|c|}\hline \\ \hline P(A) \\ \hline \end{array}, \alpha' \right), \end{array} \right) \\ \hline \end{array}$$

Let's now analyse how A and B reason about the content of (1c), given the SDRSs and models of their cognitive states that they have constructed so far.

(1) a. A: How can I get to the treasure?
    b. B: It's at the secret valley.
    c. A: But I don't know how to get there.

The DRS for (1c) is δ below (we've made the simplifying assumption that *don't know how to* is analysable as *don't know how I can*): this must be attached to β_1 with a discourse relation.[32]

[32] We have resolved the anaphor *there* to the secret valley, because this is the only candidate antecedent according to SDRT's definition of availability of antecedents, regardless of the rhetorical relation that's used to attach δ to β_1.

(δ)

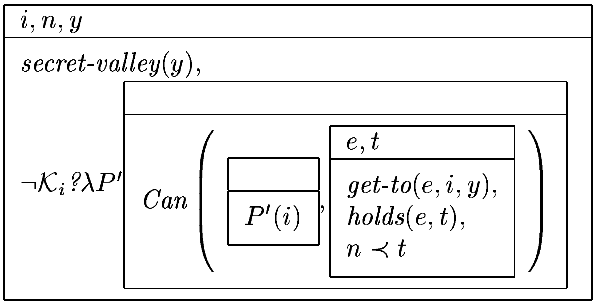

A has used the cue word *but*, which monotonically yields the relation *Contrast*(β_1, δ). This rhetorical relation, like all others, is accompanied by its own set of coherence constraints. These are defined precisely for monologue in Asher (1993). Roughly speaking, *Contrast*(α, β) in monologue requires a partial structural isomorphism and contrasting themes between α and β. These constraints capture the intuition that *Contrast*(β_1, δ) is coherent if δ indicates that some expectation arising from β_1 has been violated.[33] In the case of monologue, the expectations that arise from β_1 are essentially semantic and pragmatic entailments about the domain of discourse that β_1 refers to. But as we have seen, in the case of dialogue the expectations aren't simply recorded at the domain level. We have shown how the discourse action of uttering β_1 yields expectations about the content of the participants' cognitive states, via the axiom `Answer` (B thinks an adverbial of manner is inferrable for A from β_1); the axioms on *Can* (this adverbial of manner is the restrictor for the connective *Can*, and A must be able to do this manner for the adverbial to be an answer to the question); and the `Mutual Belief` axioms (A is aware of how B thought A should attach β_1 to α). So in order to check that δ violates an expectation arising from β_1, we must consider not only the content of β_1, but also the inferences we obtained from the utterance β_1 about the cognitive states of the participants.

We capture this by extending the extant coherence constraints in SDRT on *Contrast*. We assume that in the case of dialogue, a partial structural isomorphism and contrasting theme between δ and an expected cognitive state, given the utterance β_1, is sufficient for *Contrast*(β_1, δ) to be coherent. This is in contrast to monologue, where the partial structural isomorphism and the contrasting themes must hold between δ and β_1 itself. The rest of this section is devoted to working out the details of this suggestion.

Asher (1993) defines structural isomorphism precisely, using the notion of an embedding tree (ET). For our purposes, one can assume that an ET of a DRS is the usual DRS structure. So, for example, the ET for δ above is the following three-level tree, where $\neg K$ is the negated DRS, and β' is the DRS representing *I go to the secret valley*:

(ET$_\delta$)

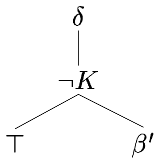

Our task is to ensure that this constituent forms the following two relations with either the constituent β_1 or with some expectation of cognitive states that was based on his uttering β_1: (a) a partial isomorphism that is sufficiently close to a bijection that it satisfies the constraints given in Asher (1993;286–291), and (b) a contrasting theme as defined in Asher (1993). In other words, we must find an SDRS or DRS which has an ET where the mapping between the nodes of the ETs is as close to a bijection as possible, and where the nodes can be assigned opposite polarities because of their semantic content (e.g., one node contains the negated information of the node it's mapped

[33] It is important to stress that this isn't a necessary condition, but intuitively it is a sufficient one.

to). The details given in Asher (1993), about what to do about when a bijection isn't possible, can be skipped over here, because in our example there is a bijection. The constituent β_1 itself doesn't meet the coherence constraints. But the part $\mathbf{K}_1$ below of A's model of B's cognitive state—which is part of KB_B—does meet the requirements:

($\mathbf{K}_1$) $$\mathcal{B}_A\left(Can\left(\boxed{P(A)}, \alpha' \right)\right)$$

It has an ET which, as shown below, forms a bijection with ET_δ (the node $\mathcal{B}_A K$ stands for the believed DRS K in $\mathbf{K}_1$, and $P(A)$ and $P'(i)$ are respectively the DRSs representing the adverbials of manner for getting to the treasure/secret valley):

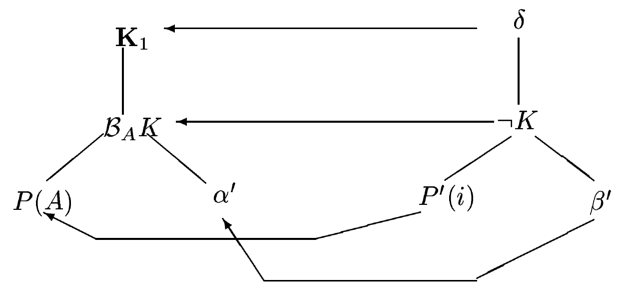

This isomorphism also yields the contrasting theme we require: δ stipulates that A can't go to the secret valley; $\mathbf{K}_1$ stipulates that A can get to the treasure; and according to the Asher (1993) definitions for assigning polarities to the nodes in the structure, this is sufficient for the nodes $\mathcal{B}_A K$ and $\neg K$ to have opposite polarities (and the node pairs α' and β', and $P(A)$ and $P'(i)$, agree in polarity). So the mapping of structures has produced the required contrasting theme, thereby ensuring that B successfully attaches δ to β_1 with *Contrast* in his SDRS.

But now, having successfully inferred $Contrast(\beta_1, \delta)$, B's model of A's cognitive state and δ contain contradictory elements. B assumed A knew how to get to the secret valley; δ says she doesn't know how she can do this. And yet δ is, nonmonotonically, part of B's model of A's beliefs via the `Sincerity` axiom given in section 4.2, which captures the cooperative principle that people believe the connections between constituents that they intend their hearers to recognise.

B was assuming that A believed he could go to the secret valley. It is for this reason that B believed that his response (1b) could be attached with *IQAP* to (1a). But the `Sincerity` axiom, when applied to δ, enables B to nonmonotonically infer that A believes he cannot. So, B must do some belief revision. A's discourse behaviour in uttering (1c) together with `Sincerity` has exposed to B that he got A's SDRS and his model of A's beliefs wrong. So, the DICE axioms have captured the following: A used B's utterance to infer what B thought A thought; detected a discrepancy; and so uttered (1c) which was specifically designed via the *Contrast* relation to enable B to detect his mistake and rectify it. Thus we have information flow from utterances to cognitive states and back again, via general axioms that encapsulate the semantics of rhetorical relations, cooperation and mutual belief.

Let's assume B drops the belief that A knows how to get to the secret valley. Then $\mathbf{K}_1$ must be removed from B's cognitive state. This means that the coherence constraints on $IQAP(\alpha, \beta_1)$ aren't met anymore. So B must drop $IQAP(\alpha, \beta_1)$ too. But now what discourse relation binds α to β_1 in B's revised SDRS? As B has inferred that A couldn't have attached β_1 to α with *IQAP*, he concludes that A must have used the axiom `NEI` to attach the constituents together with *NEI*.

So we assume he does the same.[34] Now that B's model of A's cognitive state has been revised, the *Contrast* connection between β_1 and δ isn't tenable anymore either. Again, we assume that when *Contrast* has prompted the interpreter to revise an indirect answer to something that is "not enough information", the contrast relation is also replaced with *NEI*, to indicate that responses must address the lack of A's knowledge that's asserted in δ.

Compare our analysis of (1c) with (1c′). The representation of (1c′) is δ', which has the ET shown below, where β' is again the DRS that A goes to the secret valley:

(1) c′. But I know how to get there.

(δ')

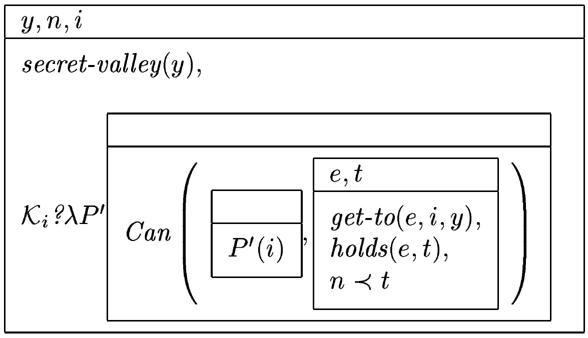

($\mathrm{ET}_{\delta'}$)

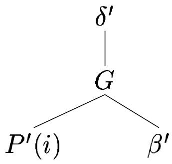

This ET fails to produce a contrasting theme with either β_1 or with the above element $\mathbf{K}_1$ of B's cognitive state. This cognitive state doesn't support the use of *but* anymore, because one can't assign opposite polarities to the node pair $\mathcal{B}_A K$ and G, as one did to $\mathcal{B}_A K$ and $\neg K$, because $\mathcal{B}_A K$ and G contain the same information. So attachment of δ' to β_1 with *Contrast* isn't possible. Hence (1c′) is incoherent in the context of (1a,b), given the way we reasoned about what B believed, from observing his discourse behaviour of uttering (1b). We have not used plans and cognitive states alone to model the difference in acceptability between (1c) and (1c′); rather, we have demonstrated the utility of folding together compositional semantics with reasoning about cognitive states, the latter being familiar from AI theories of dialogue.

As we said earlier, dialogue (2) also contains a use of contrast that exploits the information about one of the participant's cognitive states. We analyze it analogously to (1).

(2)
a. A: How do I install the assembly?
b. B: Typically, you put the assembly on before tightening the screws on the assembly.
c. A: But I tried that and it was too difficult.
d. B: So let's tighten the screws on the assembly first then.

As before, A reasons via `Answer` that B believes $IQAP(2a, 2b)$, and so A believes B thinks A can execute the plan described in (2b). A then tells B that he cannot do this plan and so sets up a contrast with B's representation of what A can do. Just as in (1), there is a conflict via `Sincerity` between B's cognitive state and the information in (2c). In the subsequent response (2d), B revises both his SDRS and his representation of A's cognitive state and then comes up

[34] We gloss over the formalisation here of how revision is pursued so as to favor achieving SDRSs between the participants that agree in as many respects as possible.

with an alternative plan to help A accomplish his goal. The discourse particle *so* indicates a result relation; the result of not being able to do the plan in β is to try the different plan in δ.

It's probably worth noting that the coherence constraints on *Contrast* that we have described here are constraints that apply in dialogue contexts generally; not just those involving *how*-questions. The difference in acceptability between (1a,b,c) and (1a,b,c′) is not peculiar to the semantics of *how*-questions, but can occur in other contexts as well, such as the *who*-question in (36):

(36)
- a. Who knows the list of applicants for the semantics job?
- b. The secretary keeps a record of all applications to the department.
- c. But I don't know where to find the secretary.
- c′. ?But I know where to find the secretary.

The compositional semantics for *who*-questions, the general axioms for reasoning with cognitive states given in section 4.2, and the coherence constraints on *Contrast* supplied here, will provide a similar account of the difference in acceptability between (36a,b,c) and (36a,b,c′) as we discussed for (1a,b,c) and (1a,b,c′).

10 Appropriacy and Contextual Relevance

Responses to questions are coherent, because the questioner can compute a rhetorical relation between these sentences and the context in his SDRS and in his model of the respondent's SDRS. Some responses, however, like (1b) and (2b), are coherent but *inappropriate*, because they do not provide the information needed for the questioner to be able to fulfil QRG. In (1) for instance, (1b) is an inappropriate response relative to A's QRG because A doesn't know how to get to the secret valley. This 'inappropriacy' was reflected in A's SDRS by the fact that β_1 attached to α with *NEI*.

It is important to stress that a response can be inappropriate in context where the relation which connects it to the interrogative is not *NEI*, but it's *QAP* or *IQAP* instead. In (1b)/(2b), the response was inappropriate for A, but B was in ostensibly no position to recognize the inappropriacy. But inappropriacy can occur for other reasons, even when A, like B, successfully attaches the response to the question with *IQAP*. For example, consider again the dialogue (16c,d) in the contexts provided in (16a) and (16b):

(16)
- a. [Context: A has just travelled by air]
- b. [Context: A has just taken a five minute taxi ride]
- c. A: Where am I?
- d. B: Helsinki.

Intuitively, (16d) is an appropriate response to (16c) in the context (16a), but not in the context (16b). Our analysis captures this intuition. Given `QRG`, A presumably wants to know a direct answer to his question that he does not already know. (16d) is a true direct answer to (16c) in both contexts. In contrast to (1), the SDRSs for A and B are the same as they both compute *QAP* between the constituents. However, what's odd about (16d) in the context of (16b) is that B should be aware that A knows what city he's in since he's only gone for a five minute taxi ride, and A assumes that B knows A knows which city he's in. So the answer B gives does not conform to what A thinks B would utter, if B was using the `Practical Syllogism` and `Cooperation`. For if B were using these axioms, then B's answer (16d) leads A to assume that B has odd assumptions about A's cognitive state, which aren't justified in this minimal context and so B is not being ideally rational. The alternative is that B is being deliberately uncooperative, which also leads us and A to wonder why. This deviation from the norm, with no explanation why, makes the dialogue sound odd.[35]

In contrast to the context (16b), it is more plausible that A doesn't know which city he's in after a long flight, and so (16d) doesn't violate the `Practical Syllogism`, `Cooperation` or `QRG` in the context of (16a). Like Ginzburg's analysis, our explanation of the difference between the two contexts makes use of some information about questioners' QRGs. But in addition, we tie inferences about questions to more general pragmatic constraints on communication, which apply not only to question answer alternations, but other forms of dialogue as well.

11 Conclusions

We have shown how links between compositional semantics at the intra-sentential level, coherence constraints on rhetorical relations at the discourse level, and intentional reasoning at the cognitive state level, contribute to the interpretation of interrogatives and responses in dialogue. A distinctive feature of our analysis was that one could compute what an utterance in dialogue meant via nonmonotonic inferences which exploited *both* compositional semantics and cognitive states. Furthermore, the way utterances are interpreted at the discourse level affected the cognitive states of the participants in the dialogue. For example, a questioner A learns what the respondent B thinks, by reasoning about the content of B's response. We modeled these cognitive effects via very general axioms which link belief, intention and desire, and which have utility beyond the dialogue situation.

These features of the analysis provided a convenient means for predicting when a contribution to a dialogue is coherent, and if coherent, when it is appropriate for achieving the participants' goals. At a general level, the axioms that we specified provide a first step towards linking formal semantic theories of intra-sentential analysis with discourse interpretation theories familiar from computational linguistics, where both rhetorical structure tools and plan recognition techniques have been used. Providing this combination in a single framework provides a richer environment for reasoning about the structure of dialogues, and their effects on the cognitive states of the participants. We demonstrated in our theory that even a small change in the surface syntax of a sentence can have a radical effect on the structure and meaning of the dialogue, in virtue of the communication links between compositional semantics at one level, and cognitive states at the other.

12 References

Nicholas Asher [1986] Belief in Discourse Representation Theory, *Journal of Philosophical Logic*, **15**, pp.127–189.

Nicholas Asher [1987] A Typology for Attitude Verbs, *Linguistics and Philosophy*, **10**, pp.125–197.

Nicholas Asher [1993] *Reference to Abstract Objects in Discourse*, Kluwer Academic Publishers.

Nicholas Asher [forthcoming] Mathematical Treatments of Discourse Contexts, in P. Dekker and M. Stokhof (eds.) *Proceedings of the Tenth Amsterdam Colloquium in Formal semantics*, pp. 21–40, ILLC publications, University of Amsterdam.

[35]The notion of appropriacy sketched here also seems to address some aspect of the problem of granularity of questions and answers. An answer to a question may not be appropriate because it is of the wrong granularity, for example, Rich Thomason (p.c.) has suggested that A told B what is in the basket if A says to B that an apple is in the basket. He hasn't told B who is in the next room when he says that a woman is next door. It seems to us that in the second case the answer may not be appropriate to B's qrg's, and that is why the answer doesn't seem helpful. Clearly, this issue requires alot more thought.

Nicholas Asher and Robert Koons [1993] The Revision of Beliefs and Intentions in a Changing World, in *Precedings of the* AAI *Spring Symposium Series: Reasoning about Mental States: Formal Theories and Applications.*

Nicholas Asher and Alex Lascarides [1994] Intentions and Information in Discourse, in *Proceedings of the 32nd Annual Meeting of the Association of Computational Linguistics (ACL94)*, pp.35–41, Las Cruces.

Nicholas Asher and Alex Lascarides [1995] Lexical Disambiguation in a Discourse Context, *Journal of Semantics*, **12**, pp.69–108.

Nicholas Asher and Pierre Sablayrolles [1995] A Typology and Discourse Semantics for Motion Verbs and Spatial PPs in French, *Journal of Semantics*, **12** 2, pp.163–210.

Nuel D. Belnap, Jr. [1985], Questions and Answers in Montague Grammar, in Stanley Peters and Esa Saarinen (eds.) *Processes, Beliefs, and Questions*, pp.165–198, D. Reidel, Dordrecht.

Michael Bennett [1979] Questions in Montague Grammar, Indiana University Linguistics Club.

Sylvain Bromberger [1962] An Approach to Explanation, in Butler, R. (ed.) *Analytical Philosophy*, Oxford University Press, pp.72–105.

Philip Cohen and Hector Levesque [1990] Rational Interaction as the Basis for Communication, in in Philip R. Cohen, Jerry Morgan and Martha E. Pollack (editors) *Intentions in Communication*, pp.221–255. Cambridge, Massachusetts: Bradford/MIT Press.

Ann Copestake (1992) *The Representation of Lexical Semantic Information*, PhD thesis, University of Sussex, Cognitive Science Research Paper CSRP–280.

Mary Dalrymple, Makoto Kanazawa, Yookyung Kim, Sam Mchombo, and Stanley Peters (in press) Reciprocal Expressions, to appear in *Linguistics and Philosophy.*

Cynthia Dwork and Yarum Moses [1986] Knowledge and Common Knowledge in a Byzantine Environment I: Crash Failures, in J. Halpern (ed.) *Theoretical Aspects of Reasoning About Knowledge: Proceedings of the 1986 Conference*, pp149–169.

Tim Fernando [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.

Jonathan Ginzburg [1993] Resolving Questions, technical report HCRC/RP-46, Human Communication Research Centre, University of Edinburgh.

Jonathan Ginzburg [1995] Resolving Questions part I, *Linguistics and Philosophy*, **18 (5)**, pp459–527.

Jonathan Ginzburg [1995] Resolving Questions part II, *Linguistics and Philosophy*, **18 (6)**, pp567–609.

Nancy Green and Sandra Carberry [1994] A Hybrid Reasoning Model for Indirect Answers, *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pp.58–65.

Jeroen Groenendijk and Martin Stokhof [1982] Semantic Analysis of *wh*-Complements, *Linguistics and Philosophy*, **5**.

Jeroen Groenendijk and Martin Stokhof [1984] *Studies on the Semantics of Questions and the Pragmatics of Answers*, PhD thesis, Centrale Interfaculteit, Amsterdam.

Barbara Grosz and Serit Kraus [1993] Collaborative Plans for Group Activities, in *Proceedings of the International Joint Conference in Artificial Intelligence (IJCAI93)*, Chambery, France, August 1993, pp.367–373.

Barbara Grosz and Candy Sidner [1986] Attention, Intentions and the Structure of Discourse. *Computational Linguistics*, **12**, 175–204.

Barbara Grosz and Candice Sidner [1990] Plans for Discourse, in Philip R. Cohen, Jerry Morgan and Martha A. Pollack (eds.) *Intentions in Communication*, MIT Press, pp.417–444.

Joseph Halpern and Yaram Moses [1984] Towards a Theory of Knowledge and Ignorance: Preliminary Report, in *Proceedings of the Non-Monotonic Reasoning Workshop*, AAAI Fall Symposium, October 1984, pp.125–143.

Jerry Hobbs [1985] On the Coherence and Structure of Discourse. Report No. CSLI-85-37, Center for the Study of Language and Information, October, 1985.

Jerry Hobbs, Mark Stickel, Douglas Appelt and Paul Martin [1993], Interpretation as Abduction, *Artificial Intelligence*, **63 (1–2)**, pp.69–142.

Hovy, E. H. [1988], Planning Coherent Multisentence Text, in *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics (ACL88)*, pp163–169.

Hans Kamp [1990] Propositional Attitudes, in C. Anthony Anderson, Joseph Owens (eds.), *The Role of Content in Logic, Language, and Mind*, CSLI Lecture Notes **20**, University of Chicago Press, pp.27–90.

Hans Kamp and Uwe Reyle [1993] *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*, Kluwer Academic Publishers.

Lauri Karttunen [1977] Syntax and Semantics of Questions, *Linguistics and Philosophy*, **1**.

Alex Lascarides and Nicholas Asher [1993] Temporal Interpretation, Discourse Relations and Commonsense Entailment, in *Linguistics and Philosophy*, **16**, pp.437–493.

Diane Litman and James Allen [1990] Discourse Processing and Commonsense Plans, in Philip R. Cohen, Jerry Morgan and Martha A. Pollack (eds.) *Intentions in Communication*, MIT Press, pp.365–388.

Karen Lochbaum [1993] A Collaborative Planning Approach to Discourse Understanding, Technical Report TR-20-93, Center for Research in Computing Technology, Harvard University.

Karen Lochbaum [1995] The Use of Knowledge Preconditions in Language Processing, in *Proceedings for the 14th Joint International Conference in Artificial Intelligence (IJCAI95)*, Montreal, August 1995.

William C Mann and Sandra Thompson [1986] Rhetorical Structure Theory: Description and Construction of Text, ISI technical report number RS-86-174.

Johanna Moore and Cecile Paris [1993] Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information, *Computational Linguistics*, **19** 4, pp.651–694.

Johanna Moore and Martha Pollack [1992] A Problem for RST: The Need for Multi-Level Discourse Analysis, *Computational Linguistics*, **18** 4, pp.537–544.

Ray Perrault [1990] An Application of Default Logic to Speech Act Theory, in Philip R. Cohen, Jerry Morgan and Martha A. Pollack (eds.) *Intentions in Communication*, MIT Press, pp.161-186.

Livia Polanyi [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.

Martha Pollack [1986] *Inferring Domain Plans in Question-Answering*, Doctoral dissertation, University of Pennsylvania, Philadelphia, PA.

Martha Pollack [1990] Plans as Complex Mental Attitudes, in Philip R. Cohen, Jerry Morgan and Martha A. Pollack (eds.) *Intentions in Communication*, MIT Press, pp.77–104.

Rob van der Sandt [1992] Presupposition Projection as Anaphora Resolution, *Journal of Semantics*, **9**, pp.133–377.

Searle, J. [1969] *Speech Acts*, Cambridge University Press, Cambridge.

Munindar Singh and Nicholas Asher [1993] A Logic of Intentions and Beliefs, *Journal of Philosophical Logic*, **22**, pp.513–544.

Robert Stalnaker [1978] Assertion, in Cole, Peter (ed.) *Syntax and Semantics volume 9: Pragmatics*, Academic Press.

Bonnie Webber, B. [1991] Structure and Ostension in the Interpretation of Discourse Deixis, *Language and Cognitive Processes*, Volume **6** Number 2, pp.107–135.