Pragmatics, and That's an Order

David Beaver
Stanford University

This paper concerns the semantics/pragmatics interface for natural language, and in particular the question of how presupposition accommodation and anaphora resolution should be orchestrated in a dynamic semantics. Previous dynamic systems such as DPL have relied on pre-indexation of anaphors and antecedents. It is argued that this represents a serious inadequacy. A general framework, Transition Preference Pragmatics is introduced. This framework allows semantic indeterminism to be counterbalanced by pragmatic constraints. As an example of the formal application of the framework, a dynamic system for anaphora resolution is developed. This system eliminates the need for pre-indexation by using a pragmatic module encoding a preference for inter-sentential parallelism. The result provides an extension to the empirical coverage of its dynamic predecessors, while also illustrating the generality of the new framework.

1 Introduction.

Nobody should be serious about *dynamic semantics* for natural language. Or rather, nobody should be serious in thinking that what goes under the term *dynamic semantics* is just semantics. Interesting dynamic analyses of natural language are interesting precisely because they provide formal analyses of phenomena that have previously been considered both (a) partly pragmatic, and (b) a pain in the butt. Two

[*]This chapter is based on an earlier version Beaver (1999b) distributed on cd. Like that earlier version, this chapter is respectfully dedicated to one of my partners in crime for the current volume, Johan van Benthem, who has always been a great source of inspiration. This chapter extends that earlier work in a number of ways, such as inclusion of the formalization of "Resolution Predicate Logic", which appeared first in my extended abstract for the Twelfth Amsterdam Colloquium Beaver (1999a). Thanks are due to Maria Aloni, Paul Dekker, Martina Faller and participants at the LLC workshop, and I am especially grateful to Johan van Benthem and two anonymous referees for comments on this paper.

Submitted for
Proceedings of LLC8
D. Barker-Plummer, D. Beaver, J. van Benthem, P. Scotto di Luzio (eds.)

phenomena that might leap to mind are discourse anaphora and presupposition.

I will introduce a framework for attacking such problems. It is called *Transition Preference Pragmatics*. You will soon see why. The focus of this chapter will be on a process that is a part of the phenomenon of discourse anaphora. Some would say that it is the most important part for many practical purposes. Yet it is a part that has usually been neglected in Amsterdam-style dynamic semantic analyses, presumably because it is considered either too pragmatic, too much of a pain in the butt, or both. The process is that of *anaphora resolution*, whereby the antecedent of an anaphoric expression is identified.

In section 2, I challenge some of the assumptions in current work on semantics, with respect to anaphora resolution and ambiguity resolution more generally. Then, in section 3, I outline the general framework for dealing with issues at the semantics pragmatics interface. In sections 4 and 5, some simple example applications are sketched, showing how this framework can be applied to presupposition accommodation and to a preference for non-trivial updates. Section 6 contains a more detailed example: I return to the issue of anaphora resolution, and define a formal system which uses Transition Preference Pragmatics to combine the system of DMPL Groenendijk et al. (1995) with a resolution component based on a naive form of syntactic parallelism. Finally, in section 7, I conclude by offering some problems for future investigation.

2 Little Numbers in Dynamic Semantics

We should care about anaphora resolution because it is an essential part of natural language meaning and inference.

Consider the following discourse, from Grosz et al. (1983), which is to be understood as a spoken discourse with no special emphasis of the pronouns:

(1) Susan gave Betsy a pet hamster. She reminded her that hamsters were quite shy. And then Susan left.

It follows from this discourse that Susan reminded Betsy of something. It does not follow that Betsy reminded Susan of anything. These are examples of inferences in natural language which stem from general properties of how anaphora is resolved.

I do not see these facts as being different in kind from inferences which stem from general properties of connective or operator meaning.

They are facts about what English means, and about which patterns of inference are valid in English. They are facts that a theory of the meaning of English should incorporate. But, unfortunately, they are not facts that are predicted by File Change Semantics (FCS) (Heim, 1982), Discourse Representation Theory (DRT) (Kamp and Reyle, 1993), Dynamic Predicate Logic (DPL) (Groenendijk and Stokhof, 1991b) and Dynamic Montague Grammar (DMG) Groenendijk and Stokhof (1991a).

Anaphora resolution is just a special case of ambiguity resolution. Consider the word-sense variation of "bank" in the following:

(2) a. Mary had been living it up for some time, and needed some money. She drove over to the nearest bank.

b. Mary went to a financial institution.

(3) a. Mary had been paddling upstream for sometime, and needed a rest. She steered her canoe to the nearest bank.

b. Mary went to the edge of a river.

Clearly, 2b follows from 2a (and not from 3a), while 3b follows from 3a (and not 2a). These inferences seem to be at least as robust as e.g. Aristotelian syllogisms. Perhaps more so for many speakers. It seems arguable, then, that a formal model of language meaning should capture these inferences as well as the inferences beloved of logicians, which largely center around operators and connectives. Do current formal theories of linguistic interpretation capture such inferences? In general, no. The last two decades of research on linguistic meaning has brought forth a host of dynamic theories of interpretation, such as FCS, DRT, DPL and DMG. None of these influential theories addresses the types of inferences discussed above. Nor do they even pretend to.

The level of abstraction at which much linguistic semantics is pitched would be staggering to many scholars of the textual form. They would be amazed to find out that some of the most sophisticated analyses of discourse anaphora yet proposed place almost no constraints on the resolution of pronominal expressions within a typical text beyond number/gender agreement, and simple left-right ordering of antecedents and anaphors. The only substantive constraint that is placed on discourse anaphora prevents resolution to antecedents buried under operators such as negations, implications and quantifiers. But just how sophisticated has a theory to be in order to state such a constraint? And given the phenomena of modal and quantificational subordination (Roberts, 1987), should we even accept the empirical validity of that constraint?

All this is not to say that dynamic accounts of discourse anaphora have achieved nothing. Crucially, these theories not only place constraints on possible anaphoric relationships, they also provide a powerful theory of what discourses containing anaphoric expressions mean, and, in some cases, they do so without any violation of compositionality. What they do not do is say anything very surprising about the anaphora resolution process or how it is constrained.

How, then, can a dynamic theory of interpretation avoid issues of anaphora resolution and yet still lay claim to being a theory of anaphoric meaning? The key, in FCS, DPL and DMG, is a bizarre assumption, although it is an assumption that is not made in DRT. This is the assumption that anaphor-antecedent relationships are presented 'gift-wrapped' to the semantics in the form of pre-indexed NPs.[1] Who or what is supposed to do the wrapping is not made entirely clear.

Perhaps syntax does the job of pre-indexation? Yet syntax seems unlikely to do more than provide a few extra local constraints on resolution via, for example, C-command conditions.

If not syntax, then presumably pragmatics must do the job. But this would not sit well with what appears to be the architecture of DPL, DMG and FCS. For it seems that pragmatics should either operate after semantic interpretation, as is assumed e.g. in classic work on pragmatics by Grice, Austin and Searle, or at least should operate in tandem with semantics. It seems unattractive to force pragmatics to constrain semantic forms in advance of semantic interpretation. To this, defenders of DPL, DMG and FCS might counter that these models were never intended to describe a model of human sentence/discourse processing, but only of sentence/discourse meaning, so that an actual implementation of the models might somehow postpone indexation, or interleave it with semantic interpretation. However, I doubt whether the original proposers of DPL, DMG and FCS would find such a response attractive: it would certainly sit uncomfortably with the processing oriented rhetoric that one finds in the original papers introducing these theories.

In numerous articles on linguistics and philosophy, I have seen little numbers attached to words which are intended to indicate which wordsense is intended, as in

(4) Mary went to the bank$_1$.

[1] In this respect Musken's "Compositional DRT" Muskens (1996) is not like DRT but like DPL/DMG/FCS, since it does assume pre-indexation. It is possible that the account in Fernando (1999) might provide a general approach to ambiguity resolution in dynamic semantics, but I leave for later work detailed assessment of the impact of developments there on the current chapter.

To my way of thinking, word sense markers are not significantly different from anaphoric coreference indices. In both cases, it is obvious that these little numbers are purely artefactual, and not part of the real languages we wish to analyze. In both cases, semanticists routinely act as if syntax will provide semantics with pre-indexed disambiguated forms. And in neither case is this latter assumption plausible. Pragmatics must be a significant factor in identifying word senses, and yet it seems unlikely that we would want to assume that pragmatics acted entirely before semantic interpretation.

But is this a significant problem? After all decorating surface forms with numbers seems such a harmless activity. And the numbers are usually kept very tiny indeed. Yet indexation is only one way that we use to disambiguate language so as to make it more amenable to semantic analysis. In general, disambiguated levels of representation used with semantic theories may step much further from surface form. It was Bertrand Russell's thesis of misleading form that opened the flood gates (Russell, 1905). Nowadays, semanticists are happy to posit a wide range of relations between a single surface form and possibly numerous disambiguated logical forms[2]. This disambiguation accounts both for genuinely syntactic phenomena, such as phrasal attachment ambiguities, and for operator scope phenomena for which the assumption of a level of disambiguated LF is more problematic.

The assumption that semantics can operate on a disambiguated LF is not inherently problematic. It is more of a promisory note, something that must be tackled before our theory of interpretation is complete. Yet it is worrying that nearly a century after Russell first introduced semanticists to a convenient way of passing the buck, there is still remarkably little known about how disambiguation occurs.

What I find even more worrying, is that for dynamic semanticists who wish to stick by the processing rhetoric with which they have presented their theories in the past, the assumption of a disambiguated form prior to interpretation seems a simple contradiction.

DRT does not assume pre-indexation for anaphora, but instead relies upon its *resolution algorithm*. As far as anaphora resolution is concerned, then, this is the dynamic theory which comes closest to living up to the dynamic rhetoric. And Uwe Reyle's UDRT (Reyle, 1993) and its brethren incorporate the beginnings of a uniform treatment of all the different forms of ambiguity I have mentioned. I say "beginnings" because the work has concentrated more on developing strategies for avoiding ambiguity resolution than on choosing the best

[2]In Montague Grammar (Montague, 1974), this becomes a range of alternative semantically unambiguous syntactic derivations.

way of resolution when forced.

Consider anaphora resolution once more. Standard DRT places essentially similar constraints on possible anaphor-antecedent relationships to DPL, DMG and FCS. Or, rather, it places a similar lack of constraints on possible anaphor-antecedent relationships. There is no limit, for instance, on how many sentences apart a pronoun may be from its anaphor. So, while in standard DRT there is clearly a place for a more detailed account of resolution to go, it is not yet obvious just what form such an account should take.[3]

Amsterdam-style dynamic semantics has an advantage over DRT: it is less committed to any particular form of mental representation. I believe that we are at a stage where it is useful to gain a general formal perspective on ambiguity resolution, that in the long-run this will lead to important results and a clearer picture of what a dynamic model of interpretation is for. For this reason, I will present my thoughts not in DRT, but in the more abstract Amsterdam style. That is to say, I will rely on a model-theoretic characterization of information states rather than an implementation of information states using a particular style of representation. Yet I think it will be clear that the ideas I develop are in the spirit of the analysis of ambiguity resolution that has already occurred in DRT, and that the architecture of the theory mirrors parts of existing DRT treatments in many ways.

3 A strategy for pragmatics

Participants in a discourse have a model of the common ground, including both information about the world exterior to the conversation and information about the world "interior" to the conversation, i.e. information concerning what is being talked about.

In making an utterance, speakers attempt to convey ways in which they think the common ground should be modified. But hearers live in ignorance. They cannot be sure what transition the speaker intends. They do not know what ground was assumed initially common by the speaker, and, even if they had known, ambiguities of the utterance

[3] Some variants of DRT have been proposed which involve more sophisticated resolution algorithms. I am thinking here of the implementation developed within the ACORD project (Bès and Guillotin, 1992), and of (Cormack, 1993). The current paper is offered in the same spirit as these earlier proposals, albeit in a different formal setting. Roberts (1998) also discusses how a pragmatic theory might be combined with DRT. The strength of her theory, its sophistication, is also a weakness, in the sense that Roberts relies upon an incompletely specified theory of text structure and communicative intention. In that sense, the proposal in the current paper is less ambitious, but more strongly predictive.

mean that they cannot be 100% sure what the speaker takes to be the final common ground.

Here's the big picture. Or at least *a* big picture. Syntactic analysis and compositional interpretation, yield a set of alternative meanings. Each meaning is itself a set of transitions, i.e. pairs of information states conceived of as inputs and outputs, where an information state is one possible common ground. What do we need pragmatics for? The main reason we need it is to choose the right single transition, the one intended by the speaker, from amongst the set of sets of transitions provided by earlier stages of interpretation.

How then should we create a theory of pragmatics? First, we need a way to represent uncertainty about possible transitions. Second, we need to represent preferences between alternative transitions. Finally, we need to describe those preferences. The first two stages delimit the possible theories of pragmatics. The last stage is the hard one: choosing the right theory of pragmatics from amongst the set of possible theories. In fact the first two stages are so easy that I will suggest a fairly rich solution in the next section, whereas the third is so difficult that I will only make the tiniest of inroads into it in this paper.

Following Stalnaker (1970), a model of the common ground can be represented as a set of possible worlds. Relative to a world, certain constants and predicates have an interpretation in terms a domain of individuals. This domain is assumed for the purposes of this paper to be provided independently.

Note that a possible world must include not only tables and chairs, but also more abstract objects such as subjects of discussion, and facts as to who just said what to whom. For practical purposes, it is often helpful to separate out these different aspects of worlds in our models of the common ground.[4] For instance, following Heim (1982), we might use an assignment function ("sequence") for keeping track of entities that are under discussion, so that a model of the common ground (or, for Heim, a "context") becomes a set of world assignment pairs. I take it that the main reason for making such a move is convenience: a theory of discourse referents will become long-winded if it includes a complete stipulation of those properties in the world which mean that a certain object is under discussion. It is far easier to represent such facts directly. Later, I will diverge further from use of worlds alone as the basic building blocks of the theory, and include not only Heimian sequences, but also the so-called *referent systems* of Vermeulen (1994); Groenendijk et al. (1995).

[4]However, some objections to such a separation of facts about discourse from other facts about the world are given in Stalnaker (1998).

Given that W is the set of possible worlds, a model of the common ground is an element of $\mathrm{CG} = \mathcal{P}(W)$. A transition is an element of $\mathrm{T}{=}\mathrm{CG} \times \mathrm{CG}$. A meaning is a set of such transitions, thus an element of $\mathrm{M} = \mathcal{P}(\mathrm{T})$.[5] Most of this will be familiar to students of dynamic semantics.

How can uncertainty about which transition the speaker intended be modeled? Easy. If hearers are uncertain about which was intended, then they must maintain a set of alternative transitions, i.e. an object of the same type as an utterance meaning. The members of this set will be referred to as the *alternative utterance transitions* (of U, where U is some utterance).[6]

Next, a way of representing preferences about which transition was intended is needed. I will use what I take to be a non-committal approach: the preferences will be represented as a partial ordering over the elements of the set of alternative utterance transitions. The partial ordering corresponding to U, what I will refer to as the *utterance transition ordering*, will be written $\leq_U$. This partial ordering will be encoded as a set of pairs of alternative utterance transitions, i.e. a member of $\mathrm{UTO} = \mathcal{P}(\mathrm{T} \times \mathrm{T})$, modulo some further constraints such as transitivity and reflexivity.

Here, then, is the model of interpretation I wish to propose. Compositional interpretation provides a set of utterance meanings defining alternative transitions. A number of different preferential constraints together induce an utterance transition ordering. These constraints will include (1) what initial common ground is most plausibly being assumed, given previous discourse, (2) what might be useful and coherent contributions to the conversation, and (3) any other constraints on the relative plausibility of different interpretations. On the basis of the utterance transition ordering, the hearer forms a partially ordered set of output common grounds, possibly cutting highly dis-preferred transitions out of contention. The partially ordered set of output common grounds (1) models the information state of the hearer after processing the utterance, and (2) helps induce the utterance transition ordering

[5] In section 4 a partial functional space of transitions is discussed, and in section 6 the space of transitions is restricted to a purely functional space $\mathrm{CG}^{\mathrm{CG}}$. Both of these could obviously be projected back into the relational space $\mathrm{T}{=}\mathrm{CG}\times\mathrm{CG}$ which, for the sake of generality, I prefer to use in describing the Transition Preference Pragmatics framework.

[6] Note that here there is an alternative: I could have modeled not the hearer's uncertainty about the intended update, but about the intended meaning. Such uncertainty would obviously be represented as a set of meanings. The issue of whether we really need to reason about alternative transitions or about alternative meanings appears to me to be a complex one.

for the following utterance, i.e., it serves as the next input.

Doubtless there are many ways in which the model I have proposed is only approximate. There is one sense in particular in which what I have described might be thought of as *pragmatics to a first approximation*. The model reduces pragmatics to a filtering operation on alternative transitions provided by compositional interpretation. This means that the preferred common grounds after update with an utterance will only contain information that either was taken to be assumed before the utterance was made, or was compositionally derived.

A more sophisticated model would weaken the link between the compositionally derived possible output common grounds, and the common grounds resulting from the application of pragmatics. In this way, further information could be introduced.

Consider, for example, what are known to Gricean pragmaticists as *clausal* implicatures. Suppose a speaker utters a disjunction, of the form "A or B". In many contexts, if the speaker had been in possession of evidence that A was the case, then the speaker would have just said "A". Thus an implicature is triggered that the speaker does not know A to be the case, and, similarly, that the speaker does not know B to be the case. These are *clausal* implicatures, and they depend on an assumption that the speaker is complying with Grice's *Maxim of Quantity*.

Such implicatures will not be part of the information in the output of the meaning relation provided by compositional interpretation. We must conclude that a distinction is needed, to put it in traditional terms, between sentence meaning and utterance meaning. A sentence meaning is the set of transitions provided by compositional interpretation. The utterance transitions may incorporate additional pragmatic information. If a pair $\langle I, O\rangle$ is in the extension of the sentence meaning, then there will be pairs $\langle I, O'\rangle$ in the utterance meaning, where O' is any subset of O, i.e. O' contains more information than O. Furthermore, the utterance transition ordering will typically involve a preference for output states which include information contributed by a given implicature above output states that do not include this information. So, by default, speakers will prefer outputs including implicatures than those that do not. This is the property that will make (conversational) implicatures defeasible inferences, one of their defining properties.

Allowing utterance outputs to contain more information than sentence outputs might be called *pragmatics to a second approximation*. The second approximation is the assumption that utterance meanings involve communication of strictly more information than sentence meanings. However, people may speak figuratively. This means that there are few absolute constraints on the relationship between the compositionally derived sentence transition meaning and the utterance meaning. In the final analysis, almost all such constraints are preferential.

For example, if I say "The cat is on the mat", I might mean what could have been expressed literally as "Watch out, he's got a gun." I might, but I probably do not. The set of utterance transitions corresponding to an utterance of "The cat is on the mat" include transitions which take the output to states where some individual is established to have a gun, but these transitions, especially in contexts where there is a salient cat whose location is of interest, are very low on the utterance transition ordering. In some other contexts, although I find it hard to imagine exactly what they are, such bizarre transitions may be higher on the ordering.

Allowing sentence outputs to differ from utterance outputs moves us beyond the first approximation. And removing all absolute constraints on the link between sentence and utterance output takes us a level further. But I must caution that while the framework of itself is of wide applicability, applying it in full generality brings forth as many questions (e.g. about the relation between sentence and utterance meaning) as it resolves. In the meantime, there is plenty of first-approximation work to be done in Transition Preference Pragmatics. I suggest that a full treatment of implicature and figurative language, although central to the success of the enterprise, be reserved until a later date.[7]

4 Presupposition accommodation.

The perspective on pragmatics presented in this paper is a generalization of an approach that I have developed to presupposition accommodation Beaver (1995).

First, and in this I adapt from work of Stalnaker, Karttunen and Heim, the meaning of a sentence is a partial function from contexts (common grounds) to contexts. Presuppositions place constraints on the input context, and hence produce this partiality. For instance, "Mary didn't realize that her car was on fire" presupposes that Mary had a car and that it was on fire. Thus the meaning of the sentence

[7] One approach to pragmatics that may be of assistance is that in Levinson (2000). If Levinson is right, then generalized conversational implicatures can be reduced to the application of a small number of *one-size-fits-all* heuristics incorporating simple default inferences. The limited nature of the inferencing mechanism he proposes might make it easier to implement in TPP than a conservative implementation of Grice's proposals.

contains only transitions in which the input context supports these propositions.

Second, a hearer is typically uncertain about what common ground is being assumed by the speaker, although common sense induces a partial ordering over alternative common grounds. I refer to this ordering as an information ordering. For example, it is possible, although unlikely, that I or any other speaker will assume that we commonly know that Mary's car was on fire. Given that it is implausible that a speaker will make such an assumption, contexts which support that proposition will be relatively low on a hearer's information ordering. However, suppose the speaker utters

(5) Mary didn't realize that her car was on fire.

This results in a transition which maps contexts supporting the proposition that Mary's car was on fire to new contexts, and eliminates others. The resulting information ordering will only contain contexts supporting the presupposed proposition, hence yielding the effect known as *accommodation*.

There are a number of ways in which an information ordering, $\pi \in \mathcal{P}(\mathrm{CG} \times \mathrm{CG})$, can be used to induce constraints on an utterance transition ordering, $\sigma \in \mathrm{UTO}$ defined over a set of utterance transitions $\tau \in \mathcal{P}(\mathrm{CG} \times \mathrm{CG})$. This is most easily done by defining UTO according to how π ranks the inputs of transitions, ignoring completely the utterance outputs. So, we can define UTO by requiring that given any two transitions among the alternative utterance transitions, the the first transition is at least as plausible as the second transition IFF the initial common ground of the first transition is at least as plausible as the initial common ground of the second transition. Formally, for any two transitions $A = \langle A_{\mathrm{in}}, A_{\mathrm{out}} \rangle$ and $B = \langle B_{\mathrm{in}}, B_{\mathrm{out}} \rangle$, $\langle A, B \rangle \in \mathrm{UTO}$ *iff* $\langle A_{\mathrm{in}}, B_{\mathrm{in}} \rangle \in \pi$.

Using an analysis of this form, it would be possible to duplicate the main results of my analysis of presupposition accommodation in the current setting. Consider, for example, the sentence

(6) When handling a Peruvian cricket, use goggles, and be careful not to touch the red and green spiky bits.

The expression "the red and green spiky bits" presupposes the existence of such bits. But what spiky bits are these, and where are they located? Potentially, the red and green spiky bits might be a permanent feature of the world independent of any particular episode of cricket exploration. Or, these bits might be parts of Peruvian crickets. They might also belong to the goggles, or perhaps to some other entity, such as the plants on which these crickets might hypothetically live. Maybe red and green spiky bits are Peruvian cricket food?

Of all of these possibilities, the thought that comes to mind most readily when hearing the sentence is that the spiky bits are part of the cricket. In the theory I have presented elsewhere, this fact would be represented as a general (although unexplained) property of the information orderings which hearers have. This property would require that (1) there are contexts in the ordering in which every Peruvian cricket has red and green spikes, and (2) some of these contexts are ranked higher in the ordering than any contexts in which goggles have red and green spiky bits, in which the speaker's world contains some red and green spiky bits, in which Peruvian grasshopper vegetation has red and green spiky bits, or in which red and green spiky bits are Peruvian grasshopper food, etc. Given that information orderings are like this, accommodation of the proposition that Peruvian grasshoppers have red and green spiky bits would be predicted[8].

In the current setting, the information ordering would not be used directly, but instead used to rank transitions corresponding to meanings of the original sentence. In particular, the resulting ranking would need to have the property that transitions in which the input was a context in which every Peruvian cricket has red and green spikes were ranked higher than all transitions in which the input involved goggles with red and green spiky bits, and so on. The story is essentially that in Beaver (1995, 1999c), but presented from a more general perspective.

5 Combining TPP with a Static Semantics

I will now demonstrate the application of TPP to a very simple propositional example based on a static logic. In part this is to provide a warm-up exercise prior to the developments in the next section, which involve a dynamic first order logic. Yet the application described in the current section, since it hints at how a general treatment of relevance and informativity might be modeled, is also of independent interest.

Let us start with propositional logic over some alphabet of atomic proposition letters, and let a model be a set of worlds providing interpretations to those symbols such that a valuation function may be defined. The valuation function V maps a formula to a truth value

[8] Other factors are at play in (6). In particular, the coherence of the asserted conditional will be affected by what is accommodated. More generally, plausibility of the initial context is just one factor determining how presuppositions are accommodated, and the application of TPP to other factors will have to await future work.

relative to a world, i.e. $V_{M,w}(\phi) \in \{0, 1\}$ for any world w in the model M and propositional logic formula ϕ. We then let an information state (relative to a model) be a set of worlds drawn from the model, and define a hearer's update of an input state I with a single formula ϕ (the meaning of some utterance) to produce an output state O by:

$$I[\![\phi]\!]O \quad \textit{iff} \quad O = \{w \in M \mid V_{M,w}(\phi) = 1\}$$

Equally, we may define the set of possible transitions associated with ϕ as $[\![\phi]\!] = \{\langle I, O\rangle \mid I[\![\phi]\!]O\}$.

Suppose a hearer is faced with an ambiguity, and cannot be sure which of $\phi_1, \ldots, \phi_n$ the speaker meant. Then the hearer's action may be modeled as a non-deterministic update:

$$I[\![\phi_1, \ldots, \phi_n]\!]O \quad \textit{iff} \quad \exists r \leq n\ I[\![\phi_r]\!]O$$

Transition preference pragmatics can be used to reduce or eliminate such uncertainty. For instance, we may impose a preference for updates that are relevant to some prior communicative goal such as answering an open question. This suggests a promising line of research, since there are well developed accounts of pragmatic preferences across alternative responses to a question Groenendijk and Stokhof (1996). Such a dynamic pragmatic analysis of questions and answers would take us too far afield for the current work, but the basic idea can hopefully be gleaned from the formalization in TPP of a much simpler pragmatic preference, the preference for non-trivial updates.

We can define a preference ordering $>$ in terms of two constraints, the constraint that updates produce some effect, and the constraint that updates do not result in inconsistency. Consider two transitions, i.e. input/output pairs, $\langle I_1, O_1\rangle$ and $\langle I_2, O_2\rangle$. The preference is defined by:

$$\begin{array}{lll}\langle I_1, O_1\rangle > \langle I_2, O_2\rangle & \textit{iff} & I_1 \subseteq O_1 \;\&\; I_2 = O_2 \text{ OR} \\ & & O_1 \neq \emptyset \;\&\; O_2 = \emptyset\end{array}$$

Given some set of transitions T, we can easily define a set of maximal transitions *max* relative to this ordering:

$$\textit{max}(T) = \{\tau \in T \mid \neg\exists\tau' \in T\ \tau' > \tau\}$$

As I will show, the operation *max* forces (partial) resolution of ambiguity, choosing from among the alternate available transitions. Formally, we have a choice between keeping this operation at a meta-level, or introducing it into the object language. Since the latter approach will be pursued in the next section, it is helpful to see how it works. We may view the object language as incorporating a listing operation in its syntax, and can define a dynamic resolution operator "[.]" on lists of formulae. The [.] operator simply constrains updates to be maximal:

$$[\![[\phi_1, \ldots, \phi_n]]\!] \quad = \quad max([\![\phi_1, \ldots, \phi_n]\!])$$

The maximal transitions are just those which involve non-trivial update. For example, suppose a sentence was ambiguous between one meaning ϕ that was consistent with what had been previously established, encoded in the state I, and one meaning ψ that was not. Then maximal update with the meaning pair would produce exactly the same effects as ordinary update with the more acceptable meaning. That is, under these assumptions about I, ϕ, and ψ, we have that $I[\![[\phi, \psi]]\!]O$ iff $I[\![\phi]\!]O$.

This example shows how *max* constrains outputs, but in fact it may constrain inputs too. That is, a hearer may actually learn something from an utterance about what the speaker considered to be the common ground prior to the utterance. In particular, when a speaker says something with meaning ϕ, the hearer can infer (defeasibly) from the definition of maximal update that neither ϕ nor $\neg\phi$ was taken already to be in the common ground. Formally, if $I[\![[\phi]]\!]O$ then $\exists w \in I\ V_w(\phi) = 1$ and $\exists w \in I\ V_w(\phi) = 0$.[9]

This is a case, albeit a trivial one, where pragmatics truly involves a preference over transitions, and not merely a preference over states. The presupposition example in the previous section could in fact be modeled as a preference over input states, and the example to be presented in the next section uses a preference over output states. TPP covers the more general case, although even quite complex examples, such as I will now turn too, can often be modeled without recourse to the full richness of the framework.

6 Combining TPP with a Dynamic Semantics

This section has two goals. The first is to exemplify the TPP framework in the domain of anaphora resolution. The second is to relate TPP to standard accounts of dynamic semantics. To achieve these goals, I will show how a standard dynamic model of anaphora can be formally combined with a pragmatic resolution module utilizing transition preferences. I begin by introducing the standard model, a dynamic first

[9] Clearly, the definition of [.] is a little too strict to model conversation, since as soon as a speaker repeats himself, the hearer is left unable to perform an update. I leave such issues to future work.

order logic, and then present a variant in which semantics and pragmatics are integrated.

6.1 Reintroducing DMPL

The standard model I will use is Dynamic Modal Predicate Logic (DMPL) Groenendijk et al. (1995), a system which incorporates treatments of anaphora, quantification and epistemic modality. Before showing how DMPL can be integrated with a TPP pragmatics, it is first necessary to introduce DMPL. By necessity, this introduction must be rather cursory: I will not provide a detailed motivation for the semantical choices made by the original authors. I will also not discuss aspects DMPL which play no role here, such as the treatment of negation and modality. A consequence of this omission is that DMPL will seem to have an unnecessarily gothic semantics relative to the properties that are here manifested, but it is to be hoped that readers can get a general idea of what DMPL achieves, and what it does not achieve. In particular, what DMPL does not achieve is freedom from the assumption of pre-indexation.

The definitions in this section are very close to the original DMPL formulation. The basic idea of the semantics is that propositional formulae are understood as providing information updates, mapping from an input information state to an output. The states themselves carry both information about the world external to the discourse, and discourse internal information.

Discourse external information is modeled in terms of possible worlds — each state includes a subset of possible worlds, the set of all worlds compatible with information conveyed by the discourse up to that point. Models are thus pairs $\langle W, D\rangle$ of a non-empty set of worlds and a non-empty domain of individuals. A *world* is a function from individual constants to members of D, and from n-ary predicate constants to sets of n-tuples of elements of D

Discourse internal information concerns which entities are under discussion, and is modeled in terms of two sorts of functions. The first sort are the *assignment functions*, finite sequences of members of D. For an n-member sequence, n (identified with $\{0, \ldots, n-1\}$) is the *domain* of the function. For a positive integer $r \leq n$, $g(r)$ is the r-th member of the sequence. In other words, an assignment function is just a list of individuals. As we will see, a state may contain many assignment functions: in this way we encode indeterminacy as to which individuals are under discussion. Typically this indeterminacy arises from the use of an indefinite description.

The second sort of function used to model discourse internal information are the *referent systems*. Intuitively, the referents correspond directly to the meaning of referring noun phrases. The referents are named using integers, and referent systems are functions mapping referents into the domain of the assignment functions — thus referent systems are simply functions form integers to integers.

The motivation for this two tier approach combining assignment functions and referent systems, suggested in Vermeulen (1994), is discussed at length in Groenendijk et al. (1995). Essentially, the idea is that simply using assignment functions leads to logical problems when there are two quantifications over the same variable. This might occur if two indefinite noun phrases bore the same index. In dynamic semantic systems using only one level of assignment functions, e.g. Groenendijk and Stokhof (1991b), such requantification destroys information about the previous value of the variable. This in turn leads to the possibility that some transitions can be down-dates. As these earlier authors argue, this is problematic because certain notions, such as epistemic modality and entailment, are far more natural when the semantics produces a monotonic increase of information. Referent systems provide a means of separating the names of variables (referents) from the entities being named by these variables. Thus requantification over a variable need not result in any removal of pointers to entities, and the information down-date problem is avoided.

In the current proposal, referent systems will provide a natural level at which to model anaphora resolution. As we will see shortly, the system of Resolution Predicate Logic is one in which a discourse monotonically adds new discourse entities via the mechanism of assignment functions, and uses the referent systems to non-monotonically model information about which entities were mentioned by which noun phrases in the most recent sentence. Keeping information about previous referents around not only provides a clean separation of monotonic and non-monotonic components, but would also allow us to model definite descriptions referring to salient entities previously mentioned but not in the previous sentence.

We can now define the crucial notion of an *information state*: a set of *possibilities* defined relative to some model, $\langle W, D\rangle$. A possibility i is a triple $i = \langle r, g, w\rangle$ of a referent system, an assignment function and a world, such that the range of r is within the domain of g, the range of g is within the domain of the model, D, and $w \in W$. The minimal information state $\odot$ is $\emptyset \times \emptyset \times W$, where $\emptyset$ is a function with zero domain.

Next we introduce a way of talking about discourse referents. The notation for $s[x/d]$ is familiar from DMPL, with one slight modification. $s[x/d]$ means the variant of s in which the referent x is mapped to a new discourse entity, which in turn all assignment functions map to d. Whereas in DMPL the discourse entity that x is mapped to is determined in terms of the range of the referent systems in s, below I use instead the domain of the assignment functions in s.[10]

Definition D1 **(Referent setting, extension)**
Let $i = \langle r, g, w\rangle \in I$; $i' = \langle r', g', w'\rangle \in I$; n is the largest integer in the domain of g; $d \in D, s \in S$.

$$\begin{aligned} s[x/d] &= \{i[x/d] \mid i \in s\} \\ i[x/d] &= \langle r[x/(n+1)], g[(n+1)/d], w\rangle \end{aligned}$$

The semantics itself is straightforward. Below I give only the clauses for predication, equality, conjunction and existential quantification, and refer the reader to Groenendijk et al. (1995) for additional connectives and operators. Predication and equality are interpreted as filtering operations, removing possibilities incompatible with the new assertion. Conjunction is interpreted as sequencing, and existential quantification over a referent x has the effect of (i) adding a new individual to each assignment function, (ii) multiplying the number of assignment functions to model uncertainty about which individual is being mentioned, and (iii) adjusting the referent system so that x is mapped to the new member of the domain of the assignment functions.

Definition D2 **(Semantics: basic clauses)** *Let $\langle r, g, w\rangle(\alpha) = w(\alpha)$ for α an individual or predicate constant, $g(r(\alpha))$ for α a referent in the domain of r, and otherwise be undefined. Then:*

$$\begin{aligned} s[\![Rt_1 \ldots t_n]\!] &= \{i \in s \mid \langle i(t_1), \ldots i(t_n)\rangle \in i(R)\} \\ &\quad \textit{and is undefined if any component } i(\alpha) \textit{ is undefined.} \\ s[\![t_1 = t_2]\!] &= \{i \in s \mid i(t_1) = i(t_2)\} \\ s[\![\phi \wedge \psi]\!] &= s[\![\phi]\!][\![\psi]\!] \\ s[\![\exists x \phi]\!] &= \bigcup_{d \in D} (s[x/d])[\![\phi]\!] \end{aligned}$$

[10] The reason for this is that referent systems will occasionally be reset, losing information about which referents have been introduced.

6.2 Resolution Predicate Logic

I DMPL referent system, as described above, is used to prevent requantification over a variable from destroying information about the discourse entity that the variable picks out. In the system developed here, which I term Resolution Predicate Logic (RPL), this role remains, but referent systems have two additional functions. First, they represent uncertainty about which discourse entity a given syntactic variable refers to, and, second, they perform a book-keeping role, maintaining information about the grammatical role (e.g. subject) of the last NP to refer to each discourse entity. To represent uncertainty, multiple referent systems are allowed in a single information state, a possibility not found in Groenendijk et al. (1995), and which necessitates a modified notion of entailment.

The semantics for RPL is similar to that for DMPL. In particular, I have been conservative with respect to DMPL by making meanings total functions from states to states. The reason for this is pedagogical rather than theoretical, in that I am hoping to make the formal developments maximally accessible to those familiar with earlier work on dynamic semantics. It should be clear how this semantics could be *lifted* to the level of relations between states, the level which I have argued in this paper is required for a general transition preference pragmatics, and the level which used elsewhere in my treatment of presupposition (Beaver, 1995).

Some extra equipment is needed in RPL to adjust and keep track of discourse referents. The notation $s[x \downarrow n]$ is used to denote the variant of s in which all referent systems have been modified so that the referent x is mapped to the discourse entity n, which should be in the domain of assignment functions in s. The function *d-entities* maps a state to the set of integers in the domain of assignment functions in that state, so it provides a summary of all the discourse entities which have been introduced.

Definition D3 **(Revised referent setting, extension)**
Let $i = \langle r, g, w\rangle \in I$*;* $i' = \langle r', g', w'\rangle \in I$*;* n *is the largest integer in the domain of* g*;* $d \in D, s \in S$*.*

$$\begin{aligned} s[x \downarrow n] &= \{i[x \downarrow n] \mid i \in s\} \\ &\quad \textit{provided } x \in \textit{d-entities}(s)\textit{, undefined otherwise} \\ i[x \downarrow n] &= \langle r[x/n], g, w\rangle \\ \textit{d-entities}(s) &= \{n \mid \exists\langle r, g, w\rangle \in s \ \ n \in \textit{dom}(g)\} \end{aligned}$$

There are two new operators in RPL not found in DMPL: ρx introduces the referent x as a pronoun, which must be interpreted as picking out the same discourse entity as some previous referent, and $[\phi]$, to be explicated later, which is the anaphoric resolution of ϕ. The operator ρ makes an entity indeterminate between all the discourse entities which have been introduced previously:

Definition D4 **(Anaphoric operator)**

$$s[\rho x]s' \quad \textit{iff} \quad s' = \bigcup_{n \in \textit{d-entities}(s)} s[x \downarrow n]$$

The relation $\rightsquigarrow$ defined below should be read as "possibly entails", i.e. there is a way of resolving ρ-referents such that the the antecedent entails the consequent, although there may be other ways of resolving referents so that this is not the case.

Definition D5 **(Entailment)**

$$\begin{array}{rcl} \textit{resolutions}(s) & = & \{t \mid \exists r\ t = \{\langle r', g, w\rangle \in s \mid r = r'\} \neq \emptyset\}\} \\ \phi \rightsquigarrow \psi & \textit{iff} & \exists s \in \textit{resolutions}(\odot[\![\phi]\!]) \rightarrow \exists t \in \textit{resolutions}(s[\![\psi]\!])\ \ s \leq t \end{array}$$

Consider the argument:

(7) Jane likes Mary. Therefore, she likes someone.

We translate this as the following "$\rightsquigarrow$"-sound argument pattern: $\exists x\ x = j \wedge \exists y\ y = m \wedge \text{likes}(x, y) \rightsquigarrow \rho u \wedge \exists v\ \text{likes}(u, v)$

In *Transition Preference Pragmatics*, compositional derivation fixes a set of possible state transitions, and pragmatics ("to a first approximation") merely provides a preference ordering over the alternative transitions. What has been presented in this section so far, under some assumptions about translation from natural language, is a rough and ready picture of the first stage. A competent Montague Grammarian should be able to see how sentences of natural language could be translated compositionally into *sensible* RPL equivalents, and these translations lead to output states which contain multiple possible mappings of referents to discourse entities. If a single RPL output state is divided into a set of alternative possible fully resolved output states, then what we will be in a position to apply a pragmatic ordering to select from amongst those outputs. I will now show how such an ordering can be defined, and applied to RPL outputs.

Let us require that referents be integers. We will then assume a convention whereby in the translation of sentences of English to RPL, the integers are chosen according to some algorithm based on grammatical obliqueness, so that e.g. the subject of a main clause is always translated using the referent *1*.

Formulae of RPL, if they involve ρ-formulae, can lead to unresolved states, in which it is uncertain which discourse object is referred to by a given referent. The output of such a formula can be split into alternative *resolutions*, as we have seen. Suppose that for some input s, there is a set of alternative resolutions τ. How can we select the *best* member of τ? That is the job which a pragmatic theory must be able to answer. But here, rather than delving into the question of which is the right pragmatic theory, I propose a token theory to show how in principle the job can be done.

The token theory can be stated as follows: (1) transitions are ranked according to the degree of parallelism between input and output; (2) parallelism is measured by the number of referents in the output which are mapped onto the same discourse object as the corresponding referent in the input; (3) a preference will be assumed for parallelism of referents introduced by the syntactically least oblique NPs. Thus we will prefer transitions that maintain subject parallelism, but if that fails to decide, we will prefer transitions providing direct-object parallelism, and so on. This is the basis of the *max* predicate, introduced below: $max_s(\tau)$ picks out the subset of resolutions in τ which are most strongly parallel with s. The definition makes use of a partial function $\mathcal{R}$ which maps a state involving only one referent system onto that referent system.

The resolution preferences to be used embody a naive form of parallelism, although my intention is to display how, in principle, any of a large class of theories of anaphora resolution would fit into the framework of transition preference pragmatics.[11] The result dispenses with the indexing assumption, so that indexation can no longer be seen as a point dividing DRT from other dynamic systems. In fact RPL provides a model of resolution empirically superior to standard DRT, since RPL incorporates pragmatic constraints on resolution which are not found in standard DRT.

Definition D6 **(Maximal Transitions)**

$$\begin{array}{rcl} \mathcal{R}(s) & = & r \textit{ if } s \neq \emptyset \textit{ and } \forall \langle r', g, w\rangle \in s\ r = r' \\ & & \textit{else undefined} \\ max_s(\tau) & = & \{t \in \tau \mid \neg\exists u \in \tau\ \exists n\ \ \mathcal{R}(s)(n) = \mathcal{R}(u)(n) \neq \mathcal{R}(t)(n)\ \& \end{array}$$

[11] In a longer version of the paper, available on request, I show how the Centering model Grosz et al. (1995) can be substituted.

$$\forall n' < n \;\; (\mathcal{R}(s)(n') = \mathcal{R}(t)(n') \;\Leftrightarrow\; \mathcal{R}(s)(n') = \mathcal{R}(u)(n'))\}$$

I now define a resolution operator [.] which utilizes the definition of maximal transitions. The input, s, of $[\phi]$ should be the output of a previous formula. A variant of s with all component referent systems emptied of information (the effect of "*clean*") is updated with ϕ to produce an intermediary state which may involve unresolvedness. The predicate *resolutions* breaks this state up into alternative fully resolved outputs, and the predicate *max* selects between these. To allow for the possibility of multiple maximal states, the union of maximal states is taken to produce the final output. In most cases this will be a fully resolved state, and will be the output which involves maximal parallelism with the input.

Definition D7 **(Resolution Operator)**

$$\begin{array}{rcl} \mathit{clean}(s) & = & \{\langle \emptyset, g, w\rangle \mid \exists r \langle r, g, w\rangle \in s\} \\ s[\![[\phi]]\!] & = & \cup \mathit{max}_s(\mathit{resolutions}(\mathit{clean}(s)[\![\phi]\!])) \end{array}$$

In the example below, the parallelism constraint favors outputs of the second sentence of (8) which map referent 1 onto the same discourse entity as in the outputs of the first sentence. It is thus an easily demonstrated formal property of the system defined that the translation of the discourse in (8) entails that of the sentence in (9).

(8) a. A soldier meets a sailor. She likes her.

b. $\exists 1\,\mathrm{soldier}(1) \;\wedge\; \exists 2\,\mathrm{sailor}(2) \;\wedge\; \mathrm{meets}(1,2)$
$\wedge\; [\rho 1 \wedge \rho 2 \wedge \mathrm{likes}(1,2)]$

(9) a. A soldier likes a sailor.

b. $\exists 1\,\mathrm{soldier}(1) \;\wedge\; \exists 2\,\mathrm{sailor}(2) \;\wedge\; \mathrm{likes}(1,2)$

RPL is a methodological stepping stone on the way to a fragment of natural language. The latter could be achieved using an embedding into type theory, following the approach pioneered in Muskens (1996), although this goes beyond the scope of the current paper.

The example of TPP that I have detailed, RPL, will make correct predictions about anaphora only to the extent which the parallelism analysis itself makes correct predictions. But in principle, any theory of resolution which could be expressed as a preference ordering over alternative transitions could be used instead of parallelism, including theories which defined only a partial ordering, and thus left pronoun reference underspecified.

7 Discussion

This paper is programmatic. I have suggested a level at which a wide variety of pragmatic constraints may be stated. But I have not said how we are to combine constraints from differing sources. This is the really important question to answer in future work. In the remainder of this section, I will hint at the scale of the problem.

Suppose that a particular constraint on preferred discourse relation clashes with a constraint derived from our common-sense knowledge on what is a plausible fact about the world. For instance, consider the following discourse segment:

(10) Mary went for a walk with her puppy, Claudia. She had a bone in her mouth.

Parallelism predicts that Mary had a bone in her mouth whereas common sense tells us it was Claudia[12].

This might lead us to hypothesize that hearers always prefer the meaning which is *a priori* most probable, regardless of what consequences this might have for text coherence. But such a hypothesis cannot be correct. Consider, for example, the following discourse segment:

(11) Mary recently went water-skiing with her mother, and is a keen follower of professional wrestling. She is 84 years old.

Clearly, common sense knowledge of the world makes it *a priori* more likely that Mary's mother is 84 years old than that Mary is. However, the text is most naturally interpreted to mean that Mary is 84 years old, in spite of this having a somewhat jarring effect on the reader. This illustrates that sometimes text-coherence constraints win out over common sense in our interpretation of a text.

How do we know, in general, which sort of constraint is most important? Given that it is not the case that one type of constraint is always stronger than the other, should a clash of preferences between these sources result in an absence of any overall preference? Whether the task of combining constraints can be achieved within TPP, or whether it requires some modification to the basic framework, is a question I must leave to future research.

The TPP approach to dynamic ambiguity resolution is quite general, and in principle might be applied to ambiguity phenmomena other

[12]Standard centering theory, e.g. Grosz et al. (1983), would make the same prediction as parallelism here, provided, as seems plausible, the subject "Mary" was identified as the so-called *Backward Looking Center* of the first sentence.

than anaphora, in particular to scopal and lexical ambiguity. The approach I have presented has similar goals to the work on lexical disambiguation, and pragmatics more generally, by Asher and Lascarides (1994). They use a default logic to select between alternative updates in a version of DRT. What is the relation between default logics and TPP? While I have not studied this question formally, I would like to suggest the possibility that TPP be thought of as providing the foundations of a novel model theory for default logic. The logic would be interpreted in the domain of partial orderings over transitions between epsitemic states. To my knowledge, previous proposals for preference based semantics for default logics (e.g. Veltman (1996)) are based on rankings of worlds or information states, so the intrinsically richer TPP rankings (over transitions) might allow for interesting logical variants on earlier work.[13]

Finally, I would like to mention one other approach which is closely related to TPP, namely that taken in Optimality Theory (OT). In this framework, developed by Prince and Smolensky (1993), linguistic generalizations are treated as ranked default rules. Recently a number of authors Blutner (1999); Dekker and van Rooy (2000); Hendriks and de Hoop (2001); Zeevat (1999) have made proposals for a treatment of semantic and pragmatic effects in a variant of OT (bidirectional OT). In Beaver (2001) I have shown how a variant on these proposals can be applied to the principle problem discussed in the current paper, anaphora resolution. The approach I develop there is complementary to that developed in the current paper. In Beaver (2001) I concentrate on reasoning syntactically about default rules. This is the approach taken in all Optimality Theoretic work, since OT has a powerful tableau method for default reasoning, but no logic-friendly model theory. TPP, by contrast, offers an abstract model-theoretic level of description. As has been shown repeatedly in the history of logic, working in parallel on both a syntactic and a model-theoretic level of description can be fruitful, and it is my hope that these two strands will be brought together in future work.

References

Asher, N. and A. Lascarides. 1994. Lexical disambiguation in a discourse context. *Journal of Semantics*, 12:69–108.

[13]Although Asher and Lascarides use their nonmonotonic logic to reason about change, the logic itself is defined statically in terms of default preferences for propositions, not preferences for transitions.

Beaver, D. 1995. *Presupposition and Assertion in Dynamic Semantics*. PhD thesis, University of Edinburgh. To appear: Studies in Logic, Language and Information, CSLI Press.

Beaver, D. 1999a. The logic of anaphora resolution. In Dekker, P., editor, *Proceedings of the 12th Amsterdam Colloquium*, pages 61 – 66. University of Amsterdam.

Beaver, D. 1999b. Pragmatics (to a first approximation). In Gerbrandy, J., M. Marx, M. de Rijke, and Y. Venema, editors, *JFAK — Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*. Vossiuspers, Amsterdam University Press.

Beaver, D. 1999c. Presupposition: A plea for common sense. In Moss, L., J. Ginzburg, and M. de Rijke, editors, *Logic, Language and Computation*, volume 2, pages 21–44. CSLI Press.

Beaver, D. 2001. The optimization of discourse: A new perspective on centering. *Linguistics and Philosophy (to appear)*.

Bès, G. and T. Guillotin, editors. 1992. *A Natural Language and Graphics Interface: Results and Perspectives from the ACORD Project*. Springer-Verlag.

Blutner, R. 1999. Some aspects of optimality in natural language interpretation. In de Hoop, H. and H. de Swart, editors, *Papers on Optimality Theoretic Semantics*, pages 1–21. OTL, Utrecht Institute of Linguistics.

Cormack, S. 1993. *Anaphora Resolution in Discourse Representation Theory*. PhD thesis, University of Edinburgh.

Dekker, P. and R. van Rooy. 2000. Optimality Theory and Game Theory: Some Parallels. *Journal of Semantics (to appear)*.

Fernando, T. 1999. A modal logic for non-deterministic discourse processing. *Journal of Logic, Language and Information*, 8(4):445–468.

Groenendijk, J. and M. Stokhof. 1991a. Dynamic montague grammar. In Groenendijk, J., M. Stokhof, and D. Beaver, editors, DYANA *deliverable R2.2a, Quantification and Anaphora I*. Center for Cognitive Science, University of Edinburgh.

Groenendijk, J. and M. Stokhof. 1991b. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100.

Groenendijk, J. and M. Stokhof. 1996. Questions. In van Benthem, J. and A. ter Meulen, editors, *Handbook of Logic and Linguistics*. Elsevier, Amsterdam.

Groenendijk, J., M. Stokhof, and F. Veltman. 1995. Coreference and modality. In Lappin, S., editor, *Handbook of Contemporary Semantic Theory*, pages 179–213. Blackwell, Oxford.

Grosz, B., A. Joshi, and S. Weinstein. 1983. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, pages 44–49, Cambridge, Mass. Association for Computational Linguistics.

Grosz, B., A. Joshi, and S. Weinstein. 1995. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21 (2):203–226.

Heim, I. 1982. *On the semantics of Definite and Indefinite Noun Phrases*. PhD thesis, Umass. Amherst.

Hendriks, P. and H. de Hoop. 2001. Optimality Theoretic Semantics. *Linguistics and Philosophy*, 24(1):1–32.

Kamp, H. and U. Reyle. 1993. *From Discourse to Logic*. Kluwer.

Levinson, S. 2000. *Presumptive Meanings: the Theory of Generalized Conversational Implicature*. MIT Press, Cambridge, Mass.

Montague, R. 1974. *Formal Philosophy*. Yale University Press.

Muskens, R. 1996. Combining montague semantics and discourse representation theory. *Linguistics and Philosophy*, 19:143–186.

Prince, A. and P. Smolensky. 1993. Optimality Theory: Constraint interaction in generative grammar. Technical report, Rutgers University Center for Cognitive Science.

Reyle, U. 1993. Dealing with ambiguities by underspecification: Construction, representation and deduction. *Journal of Semantics*, 10:123–179.

Roberts, C. 1987. *Modal Subordination, Anaphora and Distributivity*. PhD thesis, UMass. Amherst.

Roberts, C. 1998. The place of centering in a general theory of anaphora resolution. In *Centering Theory in Discourse*, pages 359–400. Clarendon Press, Oxford.

Russell, B. 1905. On denoting. *Mind*, 14:479–493.

Stalnaker, R. 1970. Pragmatics. *Synthese*, 22:272–289.

Stalnaker, R. 1998. On the representation of context. *Journal of Logic, Language and Information*, 7:3–19.

Veltman, F. 1996. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261.

Vermeulen, K. 1994. *Explorations of the Dynamic Environment*. PhD thesis, University of Utrecht.

Zeevat, H. 1999. Explaining Presupposition Triggers. In Dekker, P., editor, *Proceedings of the 12th Amsterdam Colloquium*, pages 19–24. University of Amsterdam.