

Coreference and Modality in the Context of Multi-Speaker Discourse*

Jeroen Groenendijk Martin Stokhof
Frank Veltman
ILLC/Department of Philosophy
University of Amsterdam

March 1996

1 Introduction

1.1 Context in Update Semantics

Update semantics¹ embodies a radical view on the relation between context and interpretation. The meaning of a sentence is identified with its context change potential, where contexts are identified with information states. The recursive definition of semantic interpretation is stated in terms of a process of updating an information state with a sentence. Meanings of sentences, then, are update functions. In general, these are partial functions, since the possibility to update with a sentence may depend on the fulfillment of certain constraints (presuppositions, the presence of antecedents for anaphora, etc.). It is important to note that context and interpretation are *interdependent*: the interpretation process depends on the context, and builds the context.²

The architecture of information states and the update process will depend partly on the object language, and on those features of language use

*Preparation of this paper was part of the Esprit Basic Research Project Dyana (6852).

¹See [Hei82, Hei83, Kam81] for early formal work in this tradition, which has its roots in the work of Karttunen and Stalnaker.

²Associations with the ‘hermeneutic circle’ are understandable, but perhaps are not entirely correct. After all, there is a difference between a circle and a spiral.

that are taken into consideration. As in [GSV96], our object language in this paper is the language of modal predicate logic, enriched with demonstratives. Contextual features that are taken into account are restricted to anaphoric relations between variables and quantifiers, and the role of epistemic modalities. In the present paper we turn to information exchange in multi-speaker discourse.³ Our main aim is to shed some light on anaphoric relations across utterances of different speakers, and on the role of epistemic modalities in information exchange.

1.2 Two Kinds of Information

In [GSV96] a fundamental distinction is made between two kinds of information: information about the world, and discourse information. In general, discourse information concerns matters which are relevant to the linguistic interpretation process as such. It is subservient to the primary goal of gathering information about the world.⁴

Discourse information is typically of a temporary nature: once (part of) the discourse is closed off, it can be discarded. And, under the assumption that all participants in a conversation have equal, and full, access to which utterances are made and by whom, discourse information is typically shared knowledge that cannot fail to be part of the common ground. Since the object language is a logical language, and assuming that the participants conversing in it are competent speakers, there can be no doubt or disagreement about matters of discourse information.⁵

One can look upon discourse information as contextual information in the strict sense of the word: it is linguistic information about the verbal context.⁶ In principle, all information of this kind can be included in the representation of discourse information within information states.⁷ But, of

³See [FB94] for a treatment in the framework of discourse representation theory. Within dynamic semantics an early source is [Dek93], to which the present paper owes a lot. Space does not allow an extensive comparison.

⁴We restrict ourselves exclusively to informative discourse.

⁵Here the language we are concerned with is the language of modal predicate logic, so this is the language that our speech participants are supposed to ‘speak’.

⁶And, with one exception, it is the only type of contextual information that is taken into account here. The exception concerns the use of demonstratives, treated here as terms of the form *this_d*, where *d* is an actual object, viz., the object pointed at while uttering the term. In line with our general assumptions concerning discourse information, we take it that all participants in a conversation can unmistakably identify the object pointed at as such.

⁷It is not at all unlikely that, in the end, full-blown DRSS—including syntactic descrip-

course, in practice (or, rather, on the basis of methodological principles) one restricts the representation of discourse information in the information states to the bare minimum needed for the particular language and the particular features of language use that one focuses on.

In [GSV96] discourse information only serves the purpose of resolving anaphoric relations. Its representation in an information state is restricted to a so-called ‘referent system’.⁸ A referent system consists of a number of pegs, and a function which associates pegs with variables. Whenever a quantifier $\exists x$ is used, a new peg is introduced in the referent system, and that peg is associated with the variable x .⁹ So, discourse information is restricted to keeping track of the different things talked about (only as formal objects, i.e., as pegs), and which variables can be used to refer back to them.

An information state as a whole consists not only of discourse information (in the form of a referent system), but also of information about the world, and of a link between the two types of information. An information state here is regarded as a set of possibilities. Each possibility consists of a referent system;¹⁰ a possible world;¹¹ and an assignment function which assigns some object from the domain of that world¹² to each of the pegs present in the referent system. Information growth can take place in two ways: the referent system may be *extended* with new pegs, (re)associating variables with them and assigning them suitable objects; and/or certain possible assignments or possible worlds may be *eliminated*.¹³

tions of the sentences uttered— are the kinds of structures one needs in an appropriate representation of discourse information, as part of information states. Constructing DRSS for the sentences uttered would then correspond to building up discourse information.

⁸The notion originates with Vermeulen, see [Ver95]. Extensive discussion can be found in [Ver94, chapter 3].

⁹If the quantifier has been used before, the variable is disconnected from the peg it used to be associated with, and is now associated with the newly introduced peg instead. This means that not every peg in a referent system has to be associated with a variable at every stage. The association function is an injection.

¹⁰Within a particular state, the referent system is the same in each of the different possibilities. This reflects the fact that there can be no doubt about what discourse information there is.

¹¹Possible worlds are identified with (interpretation functions of complete) first order models.

¹²We assume that all worlds have the same domain.

¹³The definition that captures this says that a state s' is an extension of a state s iff every possibility in s' is an extension of some possibility in s . Some further terminology: If i is a possibility in s , i' in s' and i' is an extension of i , we say that i' is a descendant of i in s' . If i has one or more descendants in s' , we say that i subsists in s' . If all i in s subsist in s' we say that s subsists in s' .

1.3 Two Main Themes

In the present paper it is argued that a proper analysis of anaphoric relations across utterances of (possibly) different speakers requires a slightly richer notion of discourse information than the one just outlined above (which was developed in [GSV96]), i.e., one which takes into account who said what. More in particular, each participant has to keep track of who is responsible for the introduction of which peg. As we will illustrate in section 3, that the interpretation of an anaphoric use of a variable, and the correctness conditions for such use, differ as to whether the participant who uses the anaphor is the same as the one who introduced the peg that it is associated with.¹⁴ An account of this forms the main theme of the present paper.

Another issue which comes up almost immediately when discussing the various circumstances in which ‘cross-utterance’ anaphora occur, is that in updating their information states, the participants in the conversation have to keep track of what constitutes their *direct* information, and what is *indirect* information, obtained from ‘hearsay’ or other sources. To appreciate the difference, notice the following. In general, for an utterance to be correct, the information state of the speaker has to support the sentence being uttered.¹⁵ However, it makes a difference whether the sentence is supported solely on the basis of the direct information of the speaker, or whether indirect information plays a role, too. In the latter case, this should be explicitly indicated by modally qualifying the sentence in question.

Before turning to our actual account, we will indicate why issues concerning coreference and modality force one to look into multi-speaker discourse. In doing so, we will introduce some basic notions that play a role in interpreting such discourse.

¹⁴At this point the analysis presented here differs from the one in [FB94].

¹⁵The term ‘support’ is a technical term. Roughly speaking, an information state s supports a sentence ϕ iff s updated with ϕ does not lead to ‘real’ information growth. Each possibility in s should subsist in the resulting state. The update may lead to an increase of discourse information, though. (And, in general, it will, since the speaker now knows that she has uttered ϕ , as do the other participants, and all this is common (discourse) knowledge, about which no disagreement can arise.)

2 The Context of Discovery

2.1 Consistency and Coherence

In [GSV96] the notions of coherence and consistency are used to account for intuitions about the acceptability and unacceptability of sequences of sentences. By way of illustration we discuss some examples.

A first example is the unacceptability of (1):¹⁶

- (1) It isn't raining outside. It might be raining outside.

$$\neg p \wedge \diamond p$$

The unacceptability of (1) can be explained by its inconsistency. No information state can be updated with this sequence on pain of ending in absurdity: once an information state has been updated with $\neg p$, it is inconsistent with $\diamond p$.¹⁷

If a sequence of sentences is inconsistent, it is also incoherent, which means that no (single) information state can support it. However, incoherence does not imply inconsistency. And neither does incoherence imply unacceptability *per se*.¹⁸ A prime example is the following:

- (2) It might be raining outside. It isn't raining outside.

$$\diamond p \wedge \neg p$$

¹⁶For the moment we stick to the representation of a sequences of sentences as a conjunction, which is interpreted as sequential update of some initial information state with the conjuncts. The possibility that the conjuncts are uttered by different speakers is explicitly acknowledged.

In this context it may be worth noting that one of the central empirical claims of [GSV96] is that in case ψ is a modal statement, there is a difference between $\exists x\phi \wedge \psi$ and $\exists x(\phi \wedge \psi)$, and hence, between $\exists x\phi \rightarrow \psi$ and $\forall x(\phi \rightarrow \psi)$. Of course, a conjunction inside the scope of a quantifier can not be interpreted as a multi-speaker discourse. The present paper investigates the differences in interpretation that appear when we take the conjuncts of $\exists x\phi \wedge \psi$ to be uttered by two different speakers instead of by one and the same speaker. But for the moment, we act as if this difference does not matter. And in this section we try to choose the examples in such a way that it matters as little as possible.

¹⁷A state s is inconsistent with a formula ϕ iff updating s with ϕ leads to the absurd state. A formula is inconsistent *per se* iff there is no state with which it is consistent.

¹⁸Coherence is used here as a technical notion: a formula ϕ is coherent iff there is some (non-absurd) state s such that s supports ϕ . Incoherent sentences can never be uttered sincerely. The notion of acceptability is used as a pre-theoretic notion, pertaining to linguistic intuitions.

This sequence of sentences is consistent. At the same time, no *single* information state can support $\diamond p$ and *at the same time* support $\neg p$. The incoherence of $\diamond p \wedge \neg p$ explains its unacceptability as a single utterance by a single speaker.¹⁹ However, as a multi-speaker discourse —or as discourse in which some time elapses between the utterance of the component sentences, thus allowing for new information to come in from outside²⁰— it is acceptable. Of course, to turn this observation into a fact that is explained by a logical theory, one has to explicitly introduce multi-speaker discourse as an object of logical analysis.

The explanation of the unacceptability of $\neg p \wedge \diamond p$ in terms of its inconsistency may suggest that consistency is a necessary condition for acceptability. However, the relation between consistency and acceptability, too, is not as hard and fast as the above may have suggested. For example, the following is inconsistent, if anything is:

- (3) It is raining outside. It isn't raining outside.

$$p \wedge \neg p$$

To be sure, (3) is unacceptable, but only if it is taken to be a single-speaker utterance (and barring self-correction). As a multi-speaker discourse it has a particular feature: it is a discourse in which the second speaker contradicts the first, which may be unfortunate, but which seems insufficient reason for deeming it unacceptable. On the contrary, if a participant in a conversation finds herself in an information state which is inconsistent with a sentence uttered by another participant, it seems to be her conversational duty to report this. One way to do so is to utter a sentence that contradicts the earlier utterance. Not giving notice of dissent is a conversational vice.²¹ The other participants will wrongly assume that the participant in question can accommodate the previous utterance.

Actually, although it is consistent, something like this is also at stake in the case of our earlier example $\diamond p \wedge \neg p$ if we regard it as a two-speaker discourse. In this case, too, the second speaker ‘contradicts’ the first. If the information state of the second speaker supports $\neg p$, it cannot fail to

¹⁹See [GSV96, section 4.1] for a more elaborate discussion.

²⁰This can be indicated explicitly by putting some dots in between the two sentences in (2). The ‘multi’ in ‘multi-speaker’ in fact refers to multiple information states, be they of different speech participants at the same time, or of the same speech participant at different times. For ease of exposition the terms ‘multi-speaker’ and ‘single speaker’ will be used in this rather metaphorical way.

²¹Recall that we limit ourselves to the aspect of information exchange here.

be the case that it is inconsistent with $\diamond p$. His utterance of $\neg p$ signals this. The difference with $p \wedge \neg p$ is that in the former case the conversation may continue smoothly, at least in principle, whereas in the latter case there is real disagreement which can only be solved by further discussion, which eventually has to lead to a downgrade of at least one of the states involved if agreement is to be reached.

Now consider our first example, $\neg p \wedge \diamond p$, again. As a single speaker utterance it is obviously inconsistent, and, hence, unacceptable. But as multi-speaker discourse its unacceptability is perhaps less obvious. Confronted with the claim $\neg p$, might one not retort by pointing out that according to one's own information the possibility of p is not excluded? Notice that this case differs from the previous one: here the second speaker does not contradict the first one. There are information states which support $\diamond p$ are consistent with $\neg p$. In this case, the second speaker rather displays that she has little trust in what the first speaker said. She is non-accommodating. Again, perhaps this is to be deplored, but is it also to be ruled out as unacceptable?

The difference between this case and the other two is important. Here a rule is operative that seems to constitute a defining characteristic of the game of information exchange, viz., that one be accommodating, and update with what is said if one can, and otherwise give notice by rejecting what was said.²² It is this rule which is violated in the sequence $\neg p \wedge \diamond p$. If the information state of the second speaker supports $\diamond p$ (which it should if her utterance is to be correct in the first place), then it can either be updated with the sentence $\neg p$ uttered by the first speaker, or, if the information of the second speaker actually happens to support p itself,²³ it would also support a straightforward contradiction of what the first speaker said. So a different game is being played here, of which exchange of information about the world apparently is not the first objective.

These few observations show that the relationship between consistency, coherence and acceptability is complicated, and touches upon matters that go right to the heart of what games can be played in language. In the remainder, a provisional analysis —just scratching the surface— will be presented

²²Rules like these have been formulated and discussed in the pragmatic literature extensively, at least since the pioneering work of Stalnaker and others. The present paper has no pretense to make any systematic addition to this body of literature, and uses these ideas in a rather free way.

²³In this case $\diamond p$ would be qualitatively but not quantitatively correct. In the present paper, we leave matters of quantity aside.

of one particular type of such game: that of sincere and accommodating exchange of information about the world.

2.2 The Aloni Sequence

In this subsection we discuss somewhat extensively an example, due to Maria Aloni, which makes clear that some facts concerning the acceptability of sentences force one to take multi-speaker discourses into account.

Sofar we considered the following cases:

- Sequences of sentences which are coherent (and hence consistent). They give rise to acceptable single speaker discourses.
- Sequences of sentences which are inconsistent, and hence incoherent, and unacceptable as single speaker discourses. Some of these may be acceptable as multi-speaker discourses, but only if the inconsistency is caused by the fact that a sentence is followed by a sentence which contradicts it. Such contradictory discourses are (considered to be) acceptable multi-speaker discourses.
- Inconsistent sequences which are not contradictory. These can only be construed as a multi-speaker discourses in which one of the participants did not accommodate the utterance of one of the others, and did not give notice of this. Such non-accommodating, non-contradictory sequences are (considered to be) unacceptable, also as multi-speaker discourses.
- Sequences of sentences which are consistent, but incoherent. When viewed as a single speaker discourse such sequences are unacceptable. Under certain conditions they can be construed as acceptable multi-speaker discourses, the conditions being that each utterance in the sequence as such is coherent, and that each utterance can either be looked upon as accommodating, or as contradicting (one of) the previous utterance(s).

In the context of the present paper, the last type of sequences —consistent but incoherent— is the most interesting. Sofar we have only considered cases where it was indeed possible to interpret such a sequence as an acceptable multi-speaker discourse. Such an interpretation, however, is not always available, and the remainder of this section is devoted to an investigation of one particular example.

The following is an example of a consistent, yet incoherent sequence, which intuitively is *not* acceptable, neither as a single speaker discourse nor as a multi-speaker one:

- (4) Alfred has not done it. There is someone who might have done it. It is Alfred.

$$\neg Pa \wedge \exists x \Diamond Px \wedge (x = a)$$

It is important to note that sequence (4) as such is consistent, i.e., that there is a state s such that an update of s with (4) does not result in the absurd state. The contrary impression, that (4) is inconsistent, is perhaps due to the fact that it is so with respect to an information state in which it is known who Alfred is.²⁴ This being so, it is important to note that also in such cases (4) is not acceptable. Recall that inconsistent sequences may be acceptable when they can be interpreted as multi-speaker discourses which are contradictory. However, such an interpretation of (4) is not forthcoming, not even in situations in which it is known who Alfred is. To see that it is not, notice that the last utterance does not contradict the second one. Actually, a correct utterance of the third sentence needs a consistent update with the second sentence, otherwise the anaphor has no antecedent. And neither is it the case that the last two sentences together contradict the first one. It does *not* hold that:²⁵

$$\exists x \Diamond Px, (x = a) \models Pa$$

What does hold is that:

$$\exists x \Diamond Px, (x = a) \models \Diamond Pa$$

But that is not enough.²⁶ An information state can support $\Diamond Pa$ and at the same time be consistent with $\neg Pa$.²⁷

Summing up sofar: when it is known who Alfred is, (4) is indeed inconsistent, but it is not generally so. Inconsistent discourses can be acceptable

²⁴Technically, this means that in such a state the name a is an identifier. A name α is an identifier in a state s iff α denotes the same object in every possibility in s .

²⁵The entailment notion used here is (roughly) defined as follows: $\phi_1, \dots, \phi_n \models \psi$ iff every information state consecutively updated with $\phi_1 \dots \phi_n$ supports ψ .

²⁶If it were, then the unacceptable discourse $p \wedge \Diamond \neg p$ would come out as acceptable, too. It would then also be characterized as inconsistent but contradictory, and hence acceptable (as a two-speaker discourse).

²⁷Hence, the notion ‘ ϕ contradicts ψ ’ that is used here amounts to: every state which supports ϕ is inconsistent with ψ .

as multi-speaker speaker discourses in case the discourse is contradictory. But (4) is a non-contradictory discourse. This provides part of an explanation of the unacceptability of (4), the other part being an argument that shows that no acceptable interpretation can be found in those circumstances in which (4) *is* consistent.²⁸ That can be provided as follows.

Let us first show that indeed if (4) is uttered in a situation in which the identity of Alfred is not settled, it can be consistent. The following simple situation provides an example. Consider the state s which consists of the following two possibilities:

- The denotation of a is the object d , and the denotation of P is empty;
- The denotation of a is the object d' , and the denotation of P consists of the object d only.

We do not know who Alfred is, it is either the object d or the object d' . We know that if someone has done it, then it is the object d . And we know that if Alfred is the object d , then no-one has done it. Note that s is consistent with, and even supports, both $\neg Pa$ and $\exists x \Diamond Px$. The state s' which results from updating s with $\neg Pa \wedge \exists x \Diamond Px$ will still consist of the same two possibilities. The difference between s and s' is merely that in the latter in both possibilities (the peg with which) the variable x (is associated) is assigned the object d , since d is the only object that might have the property P .²⁹ If next we update s' with $x = a$, the possibility in which a denotes d' and in which the object d has the property P is eliminated. But the other possibility remains. Hence, (4) is consistent with s , and hence, consistent *tout court*: there exists at least one information state such that updating that state with (4) does not result in the absurd state.³⁰

This shows that, under the condition stated, (4) can indeed be a consistent sequence. Yet, we claimed, even in such circumstances it is unaccept-

²⁸That this is so is evident from the unacceptability of the following variant of (4), in which it is explicitly indicated that it is not known who Alfred is:

- (5) Alfred, whoever he is, has not done it. There is someone who might have done it.
It is Alfred.

²⁹These claims —and many more which are to follow— are based upon the recursive definition of the effects of updating information states with formulae of modal predicate logic, as provided in [GSV96].

³⁰Notice that in updating s with (4), a real extension of information about the world has taken place: it is now known who Alfred is and that no-one has done it.

able. So, the question that remains to be answered, is what explanation of its unacceptability can be given in such situations.

If we regard (4) as a single speaker discourse, it is clear what is wrong: in that case (4) is incoherent, i.e., there is no single (non-absurd) information state that supports (4). This can be seen as follows. Any state s which is to support (4) should contain a possibility i such that there is an object d in the denotation of P . For otherwise s can not support the second sentence. Furthermore, if s is to support the first sentence, a should denote some *other* object d' in that possibility i . After s has been updated with the first two sentences, this possibility i will have resulted in as many descendants as there are objects d in the denotation of P , in which x is assigned some such object d . And in each of these descendants of i , d differs from the object denoted by a . But this means that any such descendant of i will be eliminated by a further update with $x = a$. Hence, i does not subsist after an update with (4) as a whole, which means that s does not support (4). Hence, (4) is incoherent.

However, as we have seen above, the incoherence of a sequence only implies its unacceptability as a *single* speaker discourse. As such incoherence is insufficient to rule out the possibility of an acceptable *multi*-speaker interpretation. So we need to investigate this final possibility also.

Suppose some speaker utters the first sentence, $\neg Pa$. His information state should support $\neg Pa$. Furthermore, it may be assumed that other participants in the conversation are able to consistently update with $\neg Pa$. If they were not, they should react to the utterance of $\neg Pa$, and indicate their disagreement. Since, as we already saw, (4) does not contain such disagreement, it may be assumed that $\neg Pa$ is consistent with the information of all participants. Hence, each participant will actually update her information state with $\neg Pa$. As a result, all information states involved will support $\neg Pa$, after it has been uttered. Now suppose one of the participants, either the same one or a different one, utters the second sentence. By parity of reasoning, it follows that the information state of this speaker supports $\exists x \diamond Px$, and that all other participants can consistently update with it and will actually do so. The result will be that in the information state of each participant there will be some possibility i in which x is assigned some object d which belongs to the denotation of P in i , and in which the denotation of a is some object different from d . (Actually, this means that every participant is in the kind of information state described above.) But this means that none of the participants is in a position to correctly utter the last sentence $x = a$. No information state that has been consistently updated with the

first two sentences can support the last one. The possibility i that has to be present after updating with the first two sentences would not subsist after an update with the last sentence. Hence, we conclude that an acceptable multi-speaker interpretation of (4) is ruled out also. Since all possible ways of interpreting (4) lead to failure, we conclude that it is unacceptable.

This line of reasoning seems to offer a natural explanation of the unacceptability of (4). A crucial step is the assumption that if one of the participants in a conversation can not consistently update with an utterance, she should signal this by uttering a sentence which explicitly indicates that this is the case. In the absence of such a reaction it may be assumed that the participants can, and hence do, update with the sentences uttered. It is rules such as these that govern the interpretation of multi-speaker discourse.

2.3 A Non-Aloni Sequence

By way of intermezzo, we present one more example, which at first sight seems quite like the Aloni-sequence, but in fact is not:

- (6) *A*: Chris hasn't done it.
B: You [pointing at *C*] might have done it.
C: I am Chris!

- (7) $\neg Pc \wedge \diamond this_d \wedge (this_d = c)$

Its apparent likeness to the Aloni sequence notwithstanding, this sequence *is* an acceptable discourse.

To be sure, (7) is consistent, and incoherent. The former you can check for yourself by imagining that you overhear the discourse in (6). To make things easier, suppose you have the information that either *A*, or *B*, or *C* has done it. And you do not know (as apparently *B* does not either) that *C* is called Chris, although you do know that one of the three participants is. Knowing this, read the discourse. What is your information state now? You know that *A* has done it!³¹ So, you are in a proper information state, and not in the absurd state. Hence, the sequence is consistent.

That (7) is incoherent, which implies that the three utterances can not have been made by one and the same person, can be illustrated as follows:

- (8) *C*: Chris hasn't done it.
C: I might have done it.
C: I am Chris!

³¹Hint: the one who has done it, of course knows that he has. And everyone is supposed to be sincere.

Can we explain the acceptability of (6) as a multi-speaker discourse? Yes, unlike in the Aloni sequence, the last utterance does not contradict the first two taken together. Any information state that supports that *C* is called Chris, is inconsistent with the sequence consisting of the first two sentences. In the case of the Aloni sequence this is not so, if only because the last sentence contains an anaphor: it needs the one but last sentence to provide it with an antecedent.

So, the conclusion is that (7) is a consistent, incoherent, and contradictory sequence of sentences, which characterizes it as acceptable.

2.4 The Rules of the Game

The reasoning displayed above, although it did use some of the formal machinery of the semantics of dynamic modal predicate logic, was informal by common standards. To give a formally explicit account of the acceptability and unacceptability of sequences of sentences as single or multi-speaker discourses, we need to state formal definitions pertaining to the interpretation of such discourses, and define appropriate semantic notions, such as that of (multi-speaker) discourse coherence.³² Up to now, we only have a notion of single speaker coherence at our disposal.

From the discussion of the examples in the foregoing, we can extract a provisional statement of some of the rules which apparently guide the interpretation process. In the end we would want to turn these informal characterizations into formal statements, of course, but that task is beyond the scope of the present paper. In what follows we will be satisfied with taking a closer look at these matters from an informal angle.

The first rule is a speaker-oriented rule:

Rule S The information state of a speaker should justify the sentence that he utters.

This seems obvious enough. In effect, this rule is yet another formulation of the familiar Gricean maxim of Quality. The bite is, of course, in the notion of justification. We will want to explicate this informal notion using the formal concept of ‘support’. However, a simple identification of the notion of justification with the usual notion of support will not do.

³²It is such a formal notion of discourse coherence that in the end carries the burden of providing an account of intuitive acceptability judgments.

Furthermore, it needs to be explained what it means for a participant in a discourse to update his information state with an utterance of a speaker. This is captured by the following pair of hearer-oriented rules:³³

Rule H1 A participant updates her information state with each sentence which is uttered if the latter is compatible with her information state.

Rule H2 If a sentence is uttered which is incompatible with a participant's information state, then she does not update with it, but signals the incompatibility by uttering a sentence that contradicts the sentence uttered.

Again, the interesting task, at least from the point of view of formalization, is to come up with a formal explication of the informal notion of incompatibility. It is to be expected that the notion of inconsistency will play an essential role in providing such an explication. Notice that both rules presuppose that all participants hear everything that is being said, a harmless but unrealistic idealization.

Rule H1 requires participants to be *accommodating*: if they can consistently update with what is said, they do.³⁴ Rule H2 determines what happens if no consistent update is possible. If a sentence is inconsistent with an information state, this means that updating that state with that sentence would result in the absurd state. The participants avoid this by not updating at all.³⁵ However, if this situation occurs, the participant is

³³Recall the caveat made above, in footnote 22.

³⁴According to the rule a participant not only updates with the utterances of the other participants, but also with the ones she utters herself. This may seem a harmless redundancy, but it is not. By the first rule, a speaker's information state is required to support her own utterances—which means that they can not add information about the world when her information state is updated with them. So why require this update in the first place? The reason is that such an update can add discourse information. This typically happens when a speaker introduces a new discourse topic by uttering an existentially quantified sentence, such as $\exists xPx$. By updating her own information state with $\exists xPx$ the speaker guarantees not only that other participants can use the variable x to refer back to the discourse topic, but also that she can do so herself.

³⁵This is an oversimplification, of course. Any conversational move as such will lead to changes in the information states of the participants, if only to record that such a move has been made and by whom. Also, it will be noticed that the speaker's information state apparently supports the utterance in question, and so on. Such updates effects are left out of consideration in what follows, if only because in the present set-up information states are not equipped to deal with 'higher order' information, i.e., with information about the

required to give notice immediately.³⁶

Rule H2 requires an explication of the relation of contradiction between sentences in a discourse. Contrary to expectation perhaps, this is not straightforward. The qualification ‘...in a discourse’ has to be taken seriously: what is at stake is ‘contextual contradiction’. What this means is that sentences that have been uttered earlier on in the discourse (i.e., before the sentence one is supposed to contradict) may be summoned to help in interpreting an utterance as a contradiction of a preceding one.³⁷ Another requirement is that the contradiction be of a semantic nature: for any participant —no matter what state she is in— it should be clear that a certain utterance is contradicting another one. This means that we are looking for a characterization of the notion of a contradiction along the following lines:

Contradiction ϕ contradicts ψ in the context of a preceding discourse Δ iff every state state updated with Δ which justifies ϕ is incompatible with ψ

It will be clear that a formal characterization of the relation of contradiction will also rely heavily on the formalization of the concepts of justification, compatibility, and update of an information state with the utterances in a discourse.

information of other participants. Here we are only exploring the possibilities of what can already be said in a limited, ‘flat’ approach. So, when it is said that in some situation no update effects occur, what is meant is that no ‘first order’ information is updated with the informational contents of an utterance.

Note further that just ‘skipping’ utterances which are inconsistent with one’s information state means that modal statements, like the consistency test $\diamond\phi$, will never have any conversational update effects. If the test succeeds one remains in the same state, and if it does not, it is skipped, which means that one also remains in the same state. Contrary to what this may suggest, it will turn out that modal statements do play an important role in discourse.

³⁶Note that in case of straightforwardly contradicting ϕ with $\neg\phi$, the rule forces the discourse to go on and on in a childish ‘Yes! No! Yes! No!’ manner. That, when things threaten to take such a turn, one should try and resolve the disagreement, or, if that turns out to be impossible, should change the subject and continue the conversation along different lines, is left out of consideration.

³⁷A simple example is the following. Suppose a participant utters the following sequence of sentences: $\exists!xPx, \diamond(x = a)$. Suppose another participant agrees with the first utterance, but has the information that b has the property P and that $b \neq a$. He may signal the incompatibility of what is being said with his own information by uttering the following two sentences: $x \neq a, x = b$. The first of these two utterances contradicts the last utterance of the previous speaker, but needs his first utterance, if only to provide the anaphoric x with an antecedent.

A notion of discourse coherence can be defined along the following lines:

Coherent Discourse A sequence of utterances is a coherent discourse iff there exist initial information states which can be assigned to the participants in the discourse in such a way that each utterance in the sequence proceeds in accordance with the rules of the game S, H1, and H2 stated above.

The above characterizations of rules and notions provide a background for the informal explanations given earlier, such as that of the incoherence of the Aloni sequence. To turn such informal explanations into formal ones, we need to formalize these rules and notions. In particular, we have to define the following:

- justification of an utterance by an information state
- compatibility of an utterance with an information state
- contradiction of one utterance with another

Providing such definitions is not the aim of the present paper, however. Its intention is merely to provide an informal characterization of some of the notions involved.

2.5 The Aloni Sequence Revisited

Let us now consider (a slight variant of) the Aloni sequence discussed above once more. We have shown—informally—that it is an incoherent discourse: no matter how we divide the utterances between (possibly) different participants, they cannot be assigned information states in such a way that each of the utterances proceeds in accordance with the rules of conversation stated above. How would a proper formalization be able to account for this?

Different participants engaged in an informative discourse will be informed or under-informed, well-informed or ill-informed, about different things. This means that information exchange is a delicate affair. The overall goal of informative discourse is to move from a state of under-informedness to one of informedness, about a certain topic. But there is no guarantee that as the process of information exchange proceeds we will actually be well-informed, and not ill-informed. One cause of ill-informedness is that our informants happen to be ill-informed. But even if the information

we receive is correct as such, we may nevertheless become ill-informed if we draw the wrong conclusions from what we learn.

There is no general, fool-proof way to prevent this. The dangers of engaging in talk have to be accepted as a fact of life. However, there is one thing that we would want to make sure about if we can, and that is that the rules of the game themselves are stated in such a way that they prevent mismatch of information as much as possible. More in particular, we would want to avoid the rules themselves to turn out to be a source of information mismatch. For one thing, we would want to make sure that the rules are formulated in such a way that there is no mismatch between the information which justifies a speaker's utterance and the information which a hearer obtains by updating with it. A proper formalization of compatibility should prevent this: if what a hearer may learn from an utterance differs from what justified the speaker in making it, the utterance should be marked as incompatible, and a call for a reaction which signals this, should be issued. Such a formalization is not straightforward, as the following observation shows. If we would simply identify compatibility with consistency, rule H1 would require a participant to update with an utterance whenever this does not lead to the absurd state. But that will not do. To see that it does not, consider the following slight variation of the Aloni sequence discussed above.

Suppose an information state s of some speech participant supports $\neg Pa$, the first sentence in the Aloni sequence. Unlike in the original case, we assume that this time the first sentence is not uttered. Consequently, the assumption that the information states of the other participants also support $\neg Pa$, which followed in the original set-up, is no longer warranted. Only the last two sentences, $\exists x \Diamond Px$ and $x = a$ are uttered, by other participants in the discourse. Of course, there is nothing wrong with this two-sentence discourse as such. It is coherent.

Above, we gave the following simple example of an information state s which supports $\neg Pa$ and which can be consistently updated with $\exists x \Diamond Px \wedge x = a$: it consists of just two possibilities i and i' , such that $i(a) = d$ and $i(P) = \emptyset$, and $i'(a) = d'$ and $i'(P) = \{d\}$. The result of updating s with $\exists x \Diamond Px$ and $x = a$, eliminates only the possibility i' , not the possibility i . Hence, s can be consistently updated with the sequence of both sentences. Now notice that were we to identify compatibility with consistency, rule H1 would force a participant who is in state s and who is exposed to the sequence $\exists x \Diamond Px \wedge x = a$ to update with it. But this, surely, goes against intuition. Upon hearing it, a participant who is in state s should rather *refuse* to update with this sequence. For what it states is not compatible with what

she assumes she knows. Therefore, she will not update and will signify that she does not agree. She knows that something has gone wrong, that a mismatch must exist between the information of the one who uttered the last sentence and her own. In general, the exact nature of this mismatch need not be clear. In the situation at hand, all our hearer is able to infer is that either she herself is ill-informed about $\neg Pa$, or the other participant is ill-informed or underinformed. A clear symptom of information mismatch is that our hearer would not be able to ‘echo’ the utterance $x = a$ after having updated with it. The resulting information state would not justify *her* utterance of what the other participant just has said and what he himself has updated with.

These observations suggests that the notion of compatibility, far from being identical with that of consistency, requires confirmability.

3 Coreference and Modality in Discourse

3.1 Cross-Utterance Anaphora

One of the empirical motivations behind theories of dynamic interpretation is provided by the existence of anaphoric relations across sentence boundaries. Update semantics for modal predicate logic attempts to account for such relations by giving a dynamic interpretation to the existential quantifier, which allows it to bind variables outside its syntactic scope. It is worth noting that the need for such mechanisms of extended binding is even more pressing when we consider sequences of sentences uttered by different speakers. Consider the following example:

- (9) *A*: Look, a man is walking in the park.
B: Yeah, and he is wearing blue suede shoes.

There seems to be no way in which one can consider the pronoun in the sentence uttered by *B* as somehow within the scope of the indefinite in the sentence uttered by *A*. However, such sequences present no problems for the framework of update semantics. It enables us to interpret *B*’s utterance, which we can take to be simply of the form Qx , in the context of *A*’s utterance, which we assume to be represented as $\exists xPx$, allowing the quantifier to bind the variable occurring in Qx by assigning it a dynamic interpretation. To see that it does, let s_a and s_b be the initial information states of *A* and *B*. If *A*’s utterance of $\exists xPx$ is to be correct, s_a should support $\exists xPx$. Assume that s_b can be updated consistently with $\exists xPx$, and

that the resulting state supports Qx . In that case Qx is a correct utterance for B , with x bound in the required fashion. Of course, A , too, may very well be able to consistently update his state s_a with Qx , after first having updated it with his own utterance $\exists xPx$. Notice that the latter update is indeed needed: A himself, too, needs access to the discourse information provided by $\exists xPx$ in order to be able to interpret Qx with x bound by $\exists x$ in $\exists xPx$.

Of course, the fact that A and B can *interpret* each others utterance does not guarantee that *all* goes well. It is not so difficult to imagine that they assign different individuals to the variable x , a fact which may, or may not, go unnoticed, and which hence may, or may not, lead to all kinds of communicative disaster, which we need not dwell upon here.

Example (9) was carefully chosen. The opening interjections ‘Look’ and ‘Yeah’ —which do not have counterparts in the formulae $\exists xPx$ and Qx — are essential. If we leave them out, the result is less acceptable:³⁸

(10) A : A man is walking in the park.

B : He is wearing blue suede shoes.

In this exchange, B ’s utterance sounds quite odd, and the oddity clearly stems from his use of an anaphoric pronoun³⁹ to refer back to the entity introduced by A ’s use of the indefinite. It is only in particular types of circumstances that B ’s utterance is acceptable. For example, when A and B take turns in telling (or making up) a story together, this discourse is alright. But in that case, it seems, A and B are operating as a *single* agent. And, of course, if it is a single speaker that makes both utterances, the discourse is completely acceptable.

This is a remarkable fact, which certainly needs to be explained: as a sequence of utterances by a single speaker, $\exists xPx$ followed by Qx is perfectly alright, but as a sequence of two utterances by different speakers it is not

³⁸The first utterance of (10) seems more natural when it reads: ‘There is a man walking in the park’. But this is also true for (9): ‘Look, there is a man walking in the park.’, so the difference between (9) and (10) can not be explained in this way. Formally, these two variants of A ’s utterance can not be distinguished. Also, we have no clear intuition about what exactly the difference would be. This point seems to call for further investigation, but we will not pursue it in what follows.

³⁹It is the *anaphoric* nature of the expression that is relevant, not its pronominal nature: things do not improve if we replace the pronoun by an anaphoric definite description, such as ‘The man who is walking in the park’, or even ‘The man you are talking about’. Similarly, the use of the indefinite ‘Another man’, as in ‘Another man is walking in the park, too’, gives equally bad results.

acceptable, at least not without qualification. Clearly, dynamic binding is not all there is to an explanation of this type of anaphoric relations.

A further clue as to what is going on is provided by the following observation. Other circumstances in which (10) seems more or less alright is when A and B find themselves in a particular kind of observational situation.⁴⁰ There is one particular man prominently present in the visual fields of both A and B , in such a way that A can be sure that his utterance cannot fail to draw B 's attention to this individual. B realizes this, too, and it seems that it is for that reason, and no other, that he can use the anaphoric pronoun to refer to this same individual. (Note that this situation in no way presumes that there is only one man in the park, nor that A and B think there is only one.) The interjection 'Look' in (9), which is naturally thought of as being accompanied by a pointing gesture, immediately suggests this type of situation. The 'Yeah' of B explicitly indicates that he thinks he has identified the man that A intends to refer to.⁴¹

How can we model this formally? The interjection 'Look' clearly adds a deictic, demonstrative feature to A 's utterance. The closest we can get to representing this in the logical language at hand is by the following use of a demonstrative term:

$$(11) \quad \begin{array}{l} A: \exists x(x = this_d) \wedge Px \\ B: Qx \end{array}$$

If B 's initial state s_b is updated with this formula, the variable x will refer *rigidly* to the object d ,⁴² provided that it is consistent with B 's information that the object d is a man and that d walks in the park,⁴³ i.e., it is an identifier of d .⁴⁴ Per contrast, updating s_b with $\exists xPx$ need not have this effect, i.e., it does not guarantee that x will refer rigidly. Notice that the object d has to be present in the utterance situation for the demonstrative to be used in this way. Moreover, since both being a man and walking in the park are observable properties, this actually implies that B 's initial state s_b

⁴⁰It is this interpretation of the situation which is forced upon us by the interjections in (9).

⁴¹Similar observations concerning the referentiality of indefinites and pronouns in sequences like these, is can be found in [LN91], who, however, do not discuss them as multi-speaker discourses.

⁴²See [GSV96] for the relevant definition, and some discussion of the function of such expressions. In ' $this_d$ ', d is assumed to be an object available for demonstrative reference.

⁴³If the latter is not the case, B 's update with A 's utterance would lead to the absurd state, which means that B will refuse to update, and will give notice of that.

⁴⁴A term α is an identifier in a state s iff the value of α is the same in each possibility in s .

will not just be consistent with A 's utterance, but that it will even support it. As we will argue below, support is what is needed for B 's utterance being correct. If support is lacking, B 's utterance needs a modal qualification.

Examples such as (9) and (11) may suggest that it is necessary for a variable to be an identifier in an information state if a participant is to use it correctly, at least in case the variable in question was introduced by a different participant. This, however, is not so:

- (12) A : Someone has done it.
 B : He is wearing blue suede shoes.
 C : It is Alfred.

Suppose that neither A , nor B , (nor C for that matter) has any idea about who is the one who did it. Moreover, suppose that, for some reason or other, B has the information that whoever has done it wears blue suede shoes, and that C has the information that the one who has done it is Alfred. This does not imply that C knows who Alfred is. Suppose that indeed C has no idea who Alfred is. Thus she believes that Alfred has done it while at the same time she has no idea at to who has done it. Things being as they are, there is no identifier around, and this distinguishes this case from the previous one. Nevertheless, each of the three utterances individually may be acceptable in this situation.⁴⁵ The reason is that the natural interpretation of the indefinite 'someone' in A 's utterance, is that of a definite.⁴⁶ That is to say, the default translation of A 's utterance in (12) is $\exists!xPx$ rather than $\exists xPx$. The background assumption that licenses this interpretation is that 'having done it' applies to one individual only.⁴⁷

In this way, we get an intuitively appealing explanation of the difference between (10) and (12). In order for (10) to be acceptable, special circumstances need to obtain, whereas this is not necessary in case of (12). By default the indefinite in (12) is read as a definite. But this is not possible in the case of (10). The appeal to specific observational circumstances, which is

⁴⁵However, it seems that for the *sequence* of utterances in (12) to be correct, it has to be the case that not only the initial information state of A , but also those of both B and C already *support* the information that someone has done it. If not, their utterances should be modally qualified. Similarly, it should not be the case that C 's utterance, that it is Alfred who did it, is something she concludes on the basis of the information provided by B , that the suspect wears blue suede shoes. If that were the case, C 's utterance should be modally qualified, too. We will come back to this below.

⁴⁶A term α is definite in an information state s iff there are no two possibilities in s which differ only in the value assigned to α . If a term is an identifier, then it is definite, but the reverse need not hold.

⁴⁷But this is only a default assumption, of course.

what makes (10) acceptable, is a particular way of forcing the interpretation of the indefinite into that of a definite.

3.2 Coreference and Modality

However, it seems that not in all contexts definiteness is a necessary requirement for ‘cross-utterance’ anaphora. Consider the following slight variant of (10):

(13) *A*: A man is walking in the park.

B: Then he is walking his dog.

Here it does not seem necessary to assume that the pronoun is definite in *B*’s information state: it may allow for several men who are taking a stroll in the park. But it is late at night, and *B*’s information state supports the fact that any man who is out in the park at this time of day is there for no other reason than to walk his dog. No definiteness is at stake, in this case. But it is clear that it is the modal qualification in *B*’s utterance that licenses the anaphoric link in spite of the lack of definiteness. Should we leave out the ‘Then’, the result is a case similar to (10). And then we would have to imagine particular circumstances again, in which additional features, such as observational evidence, take care of the definiteness.

In the logical language at our disposal, the ‘then’ can be represented as the \Box -operator, assigning to *B*’s utterance the form $\Box Qx$. Such a formula is a test: updating a state *s* with $\Box Qx$ tests whether *s* supports *Qx*. Hence, from the point of view of a speaker, there is little difference between *Qx* and $\Box Qx$: they have the same support conditions. So why distinguish *Qx* from $\Box Qx$ in the first place? The difference shows from the hearer’s perspective.

Updating an information state *s* with *Qx* will result in the elimination of those possibilities in which the object assigned to *x* does not have the property *P*. Updating *s* with $\Box Qx$ leaves *s* as it is if in every possibility in *s* the object assigned to *x* has the property *Q*. If this condition is not met, the result would be the absurd state, which means that $\Box Qx$ is incompatible with *s*. According to conversational rule H2 no update takes place under such circumstances. This shows that, although their support conditions are the same, *Qx* and $\Box Qx$ do have different update effects. In fact, with respect to the latter one might wonder whether it is opportune to speak of updating at all: for either *Qx* is compatible with *s* and the result is *s*, or *Qx* and *s* are incompatible and no update takes place. One might conclude that in either case, nothing happens. But such a conclusion takes into account only the first part of conversational rule H2. The same rule also requires the hearer to

react, and give notice of the incompatibility. So, although $\Box Qx$ has no direct update effects on information states, it does have definite consequences for the way in which the conversation proceeds.

If someone utters $\Box Qx$ in the context of an utterance of $\exists x Px$ by another speaker, the information he provides is that $\exists x Px \rightarrow Qx$. (Which is not, of course, the information that Qx would normally provide.)⁴⁸ It seems that what happens here displays the following pattern: A utters ϕ ; B believes that $\phi \rightarrow \psi$. This means that if B can and does update with ϕ , his information state supports ψ . This suggests that B would be justified in uttering ψ , after having been informed by A that ϕ . However, this description does not quite seem to fit the facts. Consider the difference between (14) and (15):

- (14) A : John is coming to the party.
 B : Mary is coming to the party, too.

- (15) A : John is coming to the party.
 B : Then Mary is coming to the party, too.

Suppose B has no direct evidence for the fact that Mary is coming, but he does know that if John is coming, Mary is coming too. It seems that —no matter how much or how little confidence he has in the information that A provides— the right thing for B to say is what he says in (15), not what he says in (14). But why? If B is absolutely sure about his conditional information that $p \rightarrow q$, and furthermore believes that A is completely trustworthy in his utterance of p , why can he not simply say that q ?

Apparently, general rules of conversation dictate so. And for good reasons, as we shall see. In the situation at hand the justification for B 's utterance is his conditional belief, together with the information provided by A 's utterance. Conversationally speaking it is essential that this type of justification, which essentially depends on indirect evidence, be marked, so as to distinguish it from the justification B has if he has direct evidence for what he says. Let us illustrate why this is essential, by giving two simple examples.

Suppose A believes both p and $p \rightarrow \Diamond \neg q$. Now he utters p , and B answers with an utterance of q . It makes all the difference in the world for A what kind of justification B has. If B has direct evidence for q , A can, and hence will, update. The result is a state which no longer supports $p \rightarrow \Diamond \neg q$. However, if B 's justification is indirect, e.g., if it derives 'merely' from his

⁴⁸Recall the donkey equivalence: $\exists x Px \rightarrow Qx \Leftrightarrow \forall x (Px \rightarrow Qx)$.

believing $p \rightarrow q$ in combination with A 's utterance of p , then A would have to react by pointing out that his state supports $\diamond\neg q$.

This shows that the way in which the exchange proceeds is (partly) determined by the *kind of justification* participants have. Hence, distinctions, such as between direct and indirect evidence, have to be marked explicitly. And one of the roles modals play in conversation is to act as such marks.

Another example involves revision. Again, suppose A utters p , B reacts with q , and A updates with q . One reason why it is important for A to know whether or not B 's justification depends on his own utterance is that it makes a difference when, later on, A learns that $\neg p$. If B has independent, direct evidence for q , A may revise p , leaving q untouched. But if the justification of B 's q did depend on his, i.e., on A 's p , a different course has to be steered. Suppose B 's justification for q is his belief that $p \rightarrow q$ combined with A 's utterance of p . Then if A learns $\neg p$, he should question q as well.

If these observations are basically correct, as we believe them to be, definite consequences ensue for the formulation of conversational rules. A distinction has to be made between modal and non-modal statements. The utterance of a non-modal statement is justified if it is supported solely on the basis of one's own direct information, without taking into account the indirect information provided by the discourse. Modal statements, on the other hand, may be based on both types of information. This also means that the distinction between direct and indirect information has to be encoded in information states.

The distinction between direct and indirect information not only pertains to one's own information, and the information provided by others in the discourse. Consider the difference between the following two discourses:

(16) A : John is coming to the party. Mary is coming to the party, too.

(17) A : John is coming to the party. So, Mary is coming to the party, too.

The difference between (14) and (15) on the one hand, and (16) and (17) on the other, is that the latter two are single speaker discourses. The 'So' corresponds to the \square -operator. By using it, A indicates that unlike his first utterance, his second utterance is justified by conditional, i.e., indirect information. In (16) the information supporting both utterances is direct. Note that if we replace 'So' in (17) by 'Then' the result is less acceptable. Per contrast, if we replace 'Then' by 'So' in (15) we do get a correct utterance. However its meaning seems to be slightly different. It still is a sign of justification on the basis of indirect information, partly on the basis of hearsay.

But unlike the ‘Then’, the ‘So’ seems to indicate that the speaker takes it for granted that the conditional information is shared by the previous speaker.

We end this section, by returning once more to the anaphoric cases. We have seen above that if a cross-utterance anaphor, i.e., a variable, is definite in one’s own information state —as in the examples (9), (11) and (12)— it is not necessary to modally qualify one’s statement if the antecedent was introduced by another participant. But this does seem to be required if the variable is indefinite in one’s information state. However, it remains to be seen whether this is unconditionally so. The examples (14) and (15), which do not involve anaphoric links, suggest otherwise. (Cf. the remarks in footnote 45). Consider example (12) once more, in which *A* said: ‘Someone has done it’, and *B* continued with: ‘He is wearing blue suede shoes’. It seems that in this case modal qualification is not needed if, before *A*’s utterance, *B* was already aware of the fact that some definite, though as yet unidentified object has done it. However, if this is new information for *B*, a modal qualification does seem necessary, even if he did already have the conditional information that ‘If someone has done it, he is wearing blue suede shoes’. And similar remarks apply to *C*’s utterance of ‘It is Alfred’.

4 Tentative Conclusion

In the above, we have given an informal sketch of some intriguing phenomena concerning cross-utterance anaphora, and the role that the distinction between direct and indirect information seems to play in account of them. Our observations and discussions have, by and large, been rather informal. And indeed, turning them into rigorously formalized claims is a task that still awaits us. The importance of such an account would also be that it would give substance to the claim of update semantics, which is also implicit in many other theories of dynamic interpretation, that it unifies semantic and pragmatic aspects of interpretation within a single logical framework.

To which extent our discussion is particular to the kind of framework that forms its back-drop, we are reluctant to claim anything definite about. However, we do feel that the phenomena as such, though perhaps not our discussion of them, are interesting in and of themselves, and constitute a major challenge for any theory that wants to explore the ‘real life’ of multi-speaker discourse.

References

- [Dek93] Paul Dekker. *Transsentential Meditations. Ups and Downs in Dynamic Semantics*. PhD thesis, ILLC/Department of Philosophy, University of Amsterdam, Amsterdam, 1993.
- [FB94] Nissim Francez and Jonathan Berg. A multi-agent extension of DRT. Technical Report LCL 94-9, Technion - Israel Institute of Technology, Haifa, October 1994.
- [GSV96] Jeroen Groenendijk, Martin Stokhof, and Frank Veltman. Coreference and modality. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*, pages 179–213. Blackwell, Oxford, 1996.
- [Hei82] Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, 1982. Published in 1989 by Garland, New York.
- [Hei83] Irene Heim. File change semantics and the familiarity theory of definiteness. In Rainer Bäuerle, Christoph Schwarze, and Arnim von Stechow, editors, *Meaning, Use, and Interpretation of Language*. de Gruyter, Berlin, 1983.
- [Kam81] Hans Kamp. A theory of truth and semantic representation. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, editors, *Formal Methods in the Study of Language*. Mathematical Centre, Amsterdam, 1981. Reprinted in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Truth, Interpretation and Information*. Foris, Dordrecht, 1984.
- [LN91] Peter Ludlow and Stephen Neale. Indefinite descriptions: In defense of Russell. *Linguistics and Philosophy*, 14(2):171–202, 1991.
- [Ver94] Cees F.M. Vermeulen. *Explorations of the Dynamic Environment*. PhD thesis, Onderzoeksinstituut voor Taal en Spraak, Utrecht, 1994.
- [Ver95] Cees F.M. Vermeulen. Merging without mystery or: Variables in dynamic semantics. *Journal of Philosophical Logic*, 24(4):405–50, 1995.