Changing the Context. Dynamic Semantics and Discourse

Jeroen Groenendijk & Martin Stokhof
ILLC/Department of Philosophy
University of Amsterdam

This paper is an informal introduction to some aspects of dynamic semantics. It is a compilation of earlier reports on joint work with Frank Veltman. The opening section can also be found in Groenendijk et al. 1996a. Section 3 is drawn from Groenendijk et al. 1995a. Some of the discussion in section 4 derives from Groenendijk et al. 1996c.

1 Setting the stage

1.1 Context and interpretation

Within the logical-semantical tradition, the meaning of a sentence is (often) equated with its truth conditions: to know what a sentence means is to know in which circumstances it is true or false.[1] In more up-to-date approaches,[2] however, the meaning of a sentence is identified with its context change potential: to know the meaning of a sentence is to know how it changes a context.

The difference is not that the context dependent nature of interpretation is taken into account. The importance of contextual factors is generally acknowledged within traditional logical semantics, too. Usually, truth conditions are stated relative to both a model of the world, and certain other parameters which provide contextual information, such as the time and place of the utterance, its source and addressee, and possibly other features of the utterance situation.[3]

What *is* new, is the focus on context *change*: interpretation not only *depends* on the context, but also *creates* context. This is why the more fashionable approaches are

1. Formulated in terms of truth conditions, this picture seems inherently restricted to indicative sentences. However, without much difficulty it can be extended to other sentence moods. For example, in an analogous fashion, the meaning of an interrogative sentence can be equated with its answerhood conditions: to know what an interrogative sentence means is to know what under which circumstances counts as a true answer. (See Groenendijk and Stokhof 1996 for argumentation and an overview.)
2. Such as game theoretical semantics (Hintikka 1983; Hintikka and Kulas 1985), discourse representation theory (Kamp 1981; Kamp and Reyle 1993), file change semantics (Heim 1982; Heim 1983), update semantics (Veltman 1996), dynamic semantics (Groenendijk and Stokhof 1991; Chierchia 1995).
3. Within the formal semantics tradition, this development is associated with the pioneering work of Montague, Kaplan, Lewis, Cresswell. (Partee 1996 gives an extensive overview of this tradition.)

often advertised as 'dynamic'. In taking both context dependency and context change into account, dynamic approaches to interpretation confront the hermeneutic circle. Of course, it is not the observation of the interdependency of context and interpretation that is original, but rather its incorporation within a formal framework.[4]

Studying the way in which context is constructed (and deconstructed) is particularly relevant for discourse analysis. This brings another novelty to the fore. Whereas traditionally semantics concentrated on the interpretation of single sentences, dynamic theories have discovered discourse. Again, the observation that the interpretation of a sequence of sentences, more often than not, cannot simply be equated with the interpretation of the logical conjunction of its components is far from original. However, not dumping such matters in the pragmatic wastebasket, but taking them to heart in semantics proper, might be called an innovation.

1.2 Context and information

If one restricts oneself to purely informative discourse, one can look upon context change as information change, and hence upon interpretation as an incremental process of updating information. A context can be identified with an information state, and the meaning of a sentence can be characterized as an update function on information states.[5]

Information is usually partial (and need not be correct). One way to model information is to look upon an information state as a set of possibilities, viz., those possibilities which are still open according to the information. If information concerns 'the world', an information state can be identified with a set of possible worlds, each representing a different way the actual world could be as far as the information goes. On this view, extending information about the world amounts to the elimination of certain possibilities. If an information state is updated with a sentence, those worlds are eliminated in which the sentence is false, leaving only worlds in which the sentence is true.[6]

Note that dynamic interpretation is defined here in terms of truth conditions: if this would be the complete and correct picture, there would be no reason to *replace* the traditional notion of meaning as truth conditional content by the dynamic notion of information change potential. The latter notion could simply be defined on top of the former.[7]

---

4. The present paper, being of an informal nature, does not bear witness to this. But some formal background for the concepts introduced here in an informal way, can be found in Groenendijk et al. 1995b; Groenendijk et al. 1996b.
5. This view is taken, e.g., in dynamic semantics and update semantics, and in some versions of file change semantics. As will become clear shortly, discourse representation theory embodies a different perspective.
6. This so-called 'eliminative' approach to the modeling of information and information change also has a venerable ancestry, being present already in Hintikka's early work on modalities and epistemic logic.
7. This is, basically, the line pursued in early work on context change and presupposition such as that of Stalnaker, and, somewhat later, of Gazdar. For an elegant illustration of the superfluousness of the 'dynamification' of static interpretation, see Dekker 1993b, chapter 5. A thorough historical and systematic overview of the different approaches to the dynamics of interpretation, both from a linguistic and from a logical perspective, can be found in van Benthem et al. 1996.

However, there are several ways to argue that truthconditional content is not the basic notion that oils the wheels of the interpretation engine. One such way is the following.[8] Consider the contrast between the following minimal pair (due to Barbara Partee):

(1) I dropped ten marbles and found all of them, except for one. It is probably under the sofa.

(2) I dropped ten marbles and found only nine of them. ??It is probably under the sofa.

The first sentences in (1) and (2) are truthconditionally equivalent: they provide the same information about the world. Hence, if meaning is identified with truthconditional content, they have the same meaning. At the same time, however, one may observe that whereas the continuation with the second sentence in (1) is completely unproblematic, the same continuation in (2) is not equally felicitous.[9] This points towards two conclusions. The first is that, appearances notwithstanding, the opening sentences of the two examples somehow do differ in meaning, and that, hence, truthconditional content does not exhaust meaning. The second conclusion is that the update effects of a sentence are not restricted to the information about the world it conveys, but may also concern another kind of information, which has to be incorporated in the notion of an information state, too.

1.3 Information and representation

In section 1.1, dynamic theories of interpretation were characterized as subscribing to the view that meaning is context change potential. In section 1.2, it was remarked that some such theories hold that the object of change, viz., context, basically consists of information. And the marbles-example was adduced there as one kind of example that indicates that, besides information about the world, also other information has to be taken into account.

Before turning to the details of such an account, however, we want to give a very brief sketch of another view on context, which will be referred to as the 'representational view'. It localizes the dynamics of the process of interpretation in the incremental build-up of the representation of a discourse. The context of interpretation for an individual sentence is a so-called 'discourse representation structure' (DRS), a representation of the semantic content of the preceding discourse. The sentence contributes its information to the DRS that provides the context for its interpretation, by adding so-called 'discourse referents' and constraints to it.[10]

---

8. Other arguments, not involving anaphoric relations, concern presupposition, modality, conditionals and counterfactuals, defaults, tense and aspect, plurality, questions and answers. For discussion and a wealth of references, see van Benthem et al. 1996. A textbook which concentrates on the impact of dynamic semantics on empirical linguistics is Chierchia 1995.

9. Note that if there is a pause between the two utterances, then the sequence in (2) becomes just as acceptable as that in (1). The 'pragmatic effect' of the two opening sentences in all likelihood is exactly the same: we go down on our knees and help to search for the missing marble. What is remarkable, then, is that we first have to start this physical exercise to consider the second sentence in (2) felicitous, whereas in the case of (1) we also consider it so already before we start doing our gymnastics.

10. This characterization of the representational view is drawn from Kamp and Reyle 1993. In this

For example, the interpretation of the pronoun 'it' in the second sentence of the examples (1) and (2) requires that there be a suitable discourse referent in the contextual DRS to which it can be linked.[11] The opening sentence in (1) provides one. It introduces a discourse referent for the group of ten marbles which were dropped, and another discourse referent for the one among them that was not found. In the case of (2), a discourse referent for the group of ten marbles is introduced, and another one for the nine of them that were found. And although it can be *inferred* that one marble is missing, the sentence as such does not introduce a referent for it. Hence, the pronoun 'it' in the second sentence has nothing to cohere to. This is how, in principle, discourse representation theory accounts for the difference between (1) and (2).

The discourse representation structures themselves are not objects of information, but representations of information. They are of a linguistic nature, and as such are not semantic objects. Sentences and discourses are interpreted *via* an interpretation of the DRSs that represent them. The interpretation takes the form of a standard (static) truth conditional interpretation: the meaning of a DRS, and hence of the (piece of) discourse that it represents, is identified with the set of models (possible worlds) in which it is true.

The dynamics of the interpretation process resides solely in the incremental build-up, and not in the semantic interpretation, of the DRSs, and hence, of the discourses they represent. Given that the DRSs that represent them differ in form, but are true in the same models, the difference between the opening sentences of (1) and (2) is not considered to be a difference in semantic content, but one in form, unless one is prepared to look upon the representations themselves as being (parts of) the meaning. If this be the case, the assumption of a language of thought as an intermediary between language and interpretation is an essential ingredient of discourse representation theory: it counts as a mentalistic theory of meaning.

This marks the difference between a dynamic representational theory of interpretation and a dynamic semantics. In a dynamic semantics, contexts are 'objects of information', i.e., semantic objects, not linguistic ones. Consequently, what undergoes change in the dynamic process of interpretation are semantic objects, not representations. Of course, for practical purposes a dynamic semantics for a natural language might be designed using a translation procedure into a logical language. But in principle it should be possible to do without such a representational level. Hence, the resulting theory of

---

textbook on discourse representation theory, DRSs are introduced as belonging to a 'language of thought', where it is stressed that in order to play their role in a theory of meaning, the DRSs themselves are in need of (model-theoretic) semantic interpretation. Somewhat confusingly, they are sometimes referred to as information structures, a characterization which is also is used for the models in terms of which they are interpreted. Likewise, they are sometimes said to represent sentences, or larger pieces of discourse, and they are also characterized as representing the semantic content of discourse. The latter is taken here as the most appropriate description of their ontological status.

11. Discourse referents can best be compared with syntactic variables. They are expressions of the representation language. They are not themselves referents of expressions. And they (usually) do not refer to a particular object. As is generally the case with variables, their meaning resides in the variety of possible objects that can be assigned to them. For a thorough logical investigation of these matters, see Vermeulen 1994. Historically, discourse referents go back to early work of Karttunen.

meaning can remain neutral with respect to the existence and the nature of a language of thought. It is compatible with mentalism, but it is not wedded to it.[12] However, besides such abstract philosophical and methodological questions, there is also the empirical issue of descriptive adequacy: are representational and non-representational approaches equally successful in explaining the linguistic data? And, to be sure, that issue can be settled only in the long run, by detailed investigations of concrete phenomena.

2 Interlude

In the preceding section the contours of a dynamic semantics were sketched. An example was given of a difference in meaning which cannot be accounted for as a difference in truth conditions. The diagnosis was that two sentences may provide the same information about the world, but different 'discourse information'. This view was contrasted with an alternative approach which localizes the difference at a representational level, rather than at the level of semantic content.

In section 3 the potential of a dynamic semantics will be illustrated by showing that it provides a natural framework for an analysis of anaphoric definite descriptions and certain other anaphoric noun phrases in terms of contextually restricted quantification. This particular example is chosen also because it seems to present an empirical challenge for a representational approach. The discussion that follows remains at an informal level, but it takes place against the background of the more formal presentations in Groenendijk and Stokhof 1991; Groenendijk et al. 1995b; Groenendijk et al. 1996b. The various notions that are used rather casually here are intended to be in close correspondence with their formal counterparts defined in these earlier papers.

We focus on (singular) anaphoric definite descriptions, treating them—together with certain other anaphoric terms—as quantifiers, where quantification is dynamic and contextually restricted. The analysis is in line with the philosophy of Neale 1993 and Ludlow and Neale 1991, who defend a uniform Russellian, i.e., a quantificational analysis of the semantics of definites and indefinites, explaining apparent non-quantificational aspects in (epistemic) pragmatic terms. The contribution to this stock of ideas is twofold: quantification is dynamic—which accounts for binding relations across the syntactic scope of quantifiers—; and, when appropriate, restricted to context sets—which makes sense of the uniqueness preconditions of anaphoric definite descriptions and the preconditions of other kinds of anaphoric terms. The idea that (anaphoric) definite descriptions involve context dependent quantification is not new, of course. However, the mechanisms building up contextual domains have remained largely unexplicated. Dynamic semantics seems to provide a suitable framework for analyzing these mechanisms. In combination with its dynamic quantificational mechanism, it allows an easy switch between absolute and restricted quantification.

---

12. For a more extensive discussion of the issue of representationalism, and the related question of compositionality of interpretation, see Groenendijk and Stokhof 1991; Groenendijk and Stokhof 1990; Kamp 1990. Cf., also Janssen 1996.

Also, an argument will be presented against an alternative approach to anaphoric definite descriptions, which accounts for their anaphoric nature by co-indexing them with a specific term in the context. Some examples will be provided which are intended to show that—at least in some cases—co-indexing cannot do the job, whereas contextually restricted quantification can. As it seems to be the case that in those cases where co-indexing does work, contextually restricted dynamic quantification can also be used, the hypothesis that the latter is to be preferred as a general mechanism seems not unwarranted.[13]

In section 4, some attention is paid to differences in behavior of anaphoric definites in various kinds of discourse. Besides monological texts, dialogues provide another kind of context in which they may occur, with slightly different conditions on the appropriateness of their use. One of the relevant factors is the nature of the information that speech participants have at their disposal, and may or may not share. By taking a closer look at these issues, a more subtle notion of information, and information change can be obtained. However, the empirical field of definites and anaphora is vast and treacherous. Here, we can only scratch the surface, and deal with a few, relatively simple examples. Further research is called for to really take the present analysis to the test.

3 Anaphoric descriptions and context

3.1 Two kinds of information

From the discussion of the examples in (1) and (2), we concluded that information states should contain two kinds of information: information about the world, and discourse information. In the end, it is information about the world that counts, but in acquiring such information through discourse, one also has to store information pertaining to the discourse as such. For example, in order to be able to resolve anaphoric links across utterances, one has to keep track of the discourse items, viz., the 'things' which were talked about. At present, this is the only kind of discourse information we take into account.

Information about the world is modeled as a set of possible worlds. The possible worlds which are present in an agent's information state should be looked upon as alternative ways the world could be as far as the partial information of the agent goes. As information about the world grows, some such alternatives will be eliminated. According to this picture, growth of information about the world amounts to elimination of possibilities.[14]

---

13. No claim is being made that *all* anaphoric terms can be treated in this way. In earlier papers (Groenendijk and Stokhof 1991; Groenendijk and Stokhof 1990; Groenendijk et al. 1995b; Groenendijk et al. 1996b), (singular) anaphoric pronouns are analyzed by means of co-indexing, i.e., as bound variables, where the dynamics of the binding mechanism allows for variables to be bound outside the syntactic scope of a quantifier. The present paper remains neutral with respect to the question whether a co-indexing mechanism or contextually restricted quantification is most suitable for the interpretation of anaphoric pronouns.

14. According to this picture, partiality of information is modeled in terms of the presence of several

The modeling of discourse information is restricted at present to keeping track of items which are introduced by the discourse. Extending discourse information amounts to inserting new items. An initial state will contain no discourse items. As discourse goes on, the number of items grows. Once the discourse has ended, discourse information can be discarded, and the items can be deleted. Inserting and deleting items can also occur locally, triggered by the interpretation of particular parts of the discourse, even certain parts of a single sentence.

Discourse information is linked to information about the world. A link is a possible assignment of an object to each of the discourse items, an object which—relative to a particular possible world and the values of the other items—could be the value of the item in question. When a new item is added, the possible links are extended to cover the new item. More than one such extension may be possible, which means that one link can subsist in several others. It may also happen that further information provided by the discourse about the items leads to the elimination of one or more possible links. Since links are relative to possible worlds, this may lead to the elimination of a world: cut its last link and you eliminate a possible world. Discourse information can make a world of difference.

For the purpose of illustration, information states can be depicted as simple matrices, as is shown in the figures below.[15] An initial state consists of a single column, where each field in the column is filled with a possible world. The introduction of a discourse item adds a new column to the matrix.[16] The fields of the new column are filled with an object that could be the value of the item with respect to the world in the first column. Since there can be more than one such possible value, adding a new column may result in having several different rows in the new matrix, which extend the same row in the old matrix. However, an old row may also disappear, in case it is impossible to assign a suitable value for the new field with respect to that row.

---

alternatives, where these alternatives—possible worlds—are total objects. There is an obvious alternative way of picturing partiality, viz., by modeling it in terms of a partial object, a partial world or situation. According to the latter picture, growth of information amounts to extending the situation. We opt for the eliminative picture here, because it is technically more simple.

15. Pictures can be illuminating. But they can also easily mislead. Representing information states as simple matrices has its limitations. It suggests that information states are small, finite objects, whereas in fact they are usually infinite. It is also important to keep in mind that—unlike the boxes of discourse representation theory—the matrices do not represent discourse, but depict the result of interpreting discourse. They are filled with model theoretic objects, represented in the metalanguage, not with expressions of the object language.

16. We do not take into consideration here the possibility that 'discourse' items come to life by other means than explicit discourse. For example, the salient presence of an object in the visual field shared by two or more agents may lead to the creation of a discourse item, too. (Cf., footnote 9 for a case of salient absence.)

Furthermore, it may happen that, although an item is not explicitly introduced by the discourse, it is implicitly present on the basis of what has been said. The latter may be thought to occur in case of the anaphoric use of the definite 'the captain', after one has talked about a ship, without explicitly having mentioned its captain. See Dekker 1993a for an analysis of implicit arguments in a dynamic setting.

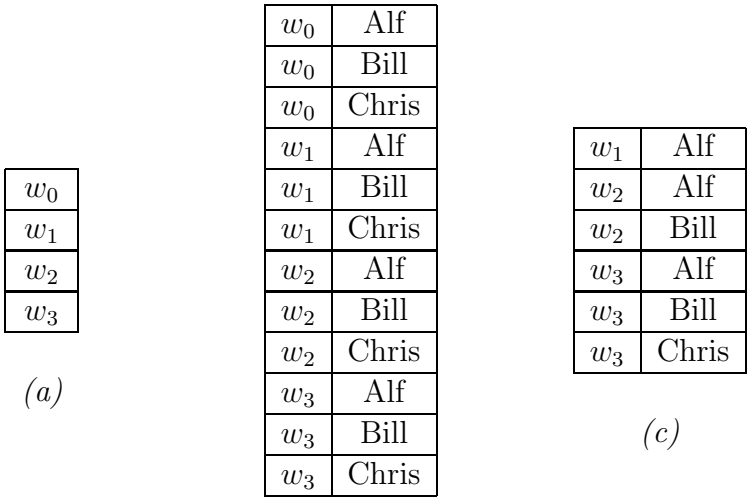

Figure 1: [Initial state] *(a)* A man *(b)* walks in the park. *(c)*

3.2 A man

Suppose an agent has the following information: Either no man walks in the park, or only Alf does, or both Alf and Bill do, or all men in the domain of discourse—Alf, Bill and Chris—are strolling there. Furthermore, he has the information that only Bill is wearing blue suede shoes.[17]

If these are the only relevant pieces of information, the information state of the agent can be depicted as in figure 1a, a one-dimensional matrix just consisting of four possible worlds. (The subscripts are used as a mnemonic device, to indicate how many men are walking in the park.)

Now suppose the agent is told the following:

(3) A man is walking in the park.

The initial information state depicted in figure 1a is transformed into state 1c, where the intermediate state 1b exemplifies the effects of processing the indefinite term 'a man'.

Interpreting an indefinite involves the introduction of a new discourse item in an information state, i.e., the addition of a new column to the matrix. With respect to each possibility in the initial state, there are three possible values to assign to the new field, since there are three men in the domain of discourse. So, for each of the four possibilities in 1a, we obtain three extensions in the intermediate state 1b, one for each man in the domain of discourse.

17. It is not thàt essential for the example, but the description of the information of the agent is to be taken in such a way that it is about objects, about the interpretations of expressions of the object language. For example, the description of the information is to be understood in such a way that the agent may very well not know which of the three men is called Alf, which one is called Bill, or which one is called Chris. In our description of the information of the agent, 'Alf', 'Bill' and 'Chris' function as expressions of the metalanguage to name these three objects. They are not the homophonous names of the language that the agent shares with other agents.

Processing the remaining predicative part of the sentence results in the elimination of rows in which the man that is the value of the new field, is not walking in the park in the world of that row. This means that in the resulting state 1c, world w_0—the world in which no man walks in the park—drops out of the picture. And each of the other three possibilities in the initial state subsists in as many extensions as there are men walking in the park in the world of that row, with one of those men as a possible value of the newly introduced discourse item.

Indefinites are interpreted in terms of dynamic existential quantification. The quantificational effect can be seen in figure 1 from the fact that world w_0, a world in which it is not the case that there is a man who walks in the park, is eliminated. This would be the only effect of ordinary 'static' existential quantification. In addition, the dynamic effect is that a new item, a new object of information, is now available in the resulting information state: a man who walks in the park. It is a partial, indefinite, non-identified object. Its presence in the information state makes it possible to refer back to hìm—the man who walks in the park.[18]

3.3 Context sets

As can be observed from the way they are depicted, information states come naturally with a contextually restricted domain of discourse. In each possibility there is not just the global domain of discourse, consisting of all the objects that live in the world of that possibility, there is also the restricted set of the objects which in that possibility are the values of the discourse items. This set is called the 'context set' of that possibility.In the states depicted in figure 2 below, the context set consists in each possibility of a single individual. And in the states depicted in figures 3b and 3c, the context set in each possibility consists of two objects.

Quantification restricted to context sets was first introduced and studied in Westerståhl 1984. He stresses the point that a context set is to be distinguished from a universe of discourse. Unlike the latter, the former is not constant over pieces of discourses. Westerståhl only considers 'the formal framework for context sets, leaving (the more difficult) question of *how* context sets are chosen to more ambitious semantic theories'. In the present set-up, context sets are not subject to choice, but are constructed (and deconstructed) in a deterministic fashion through the interpretation procedure. In principle there is a choice to be made when one meets a term in a text: that between absolute and contextually restricted quantification. But once one has opted for the latter, the relevant context sets are simply provided by the contents of the information state at that point, leaving one no further choice. The context sets do have the characteristic features of being relatively small and in constant flux, because they depend on the discourse items, which have a relatively short life span. The fact that information states come with context sets can be used to interpret anaphoric terms as contextually restricted quantifiers. The general picture is as follows.

---

18. A pioneering work on the role of information in semantics in general and on the nature of partial objects as objects of information in particular, dating from pre-dynamic days, is Landman 1986.

The update associated with an anaphoric term is characteristically partial and comes with a precondition, making a certain requirement on the actual contents of the context sets of the possibilities of the input state. Either the state has to already support the requirement, or—in case accommodation is permitted[19]—it should be consistent with it, i.e., it should be possible to update the state in such a way that afterwards it meets the requirement.[20] If the state can not (be made to) meet the precondition, the interpretation procedure aborts. If it can, the process continues along the following lines. A new discourse item is added, and the possible values of the new item are determined relative to the objects in the context sets, in a way which depends on the quantificational nature and the descriptive content of the term. Invariably, if it succeeds, the procedure as a whole will output a real extension of the input state.

3.4 The man

As for anaphoric definite descriptions,[21] they have as their precondition that within the context set of each possibility, i.e., among the values of the discourse items in a row, there is a unique object that satisfies its descriptive content. If this condition can not be fulfilled, the updating process comes to a halt. If it can, the definite description introduces a new discourse item, and in each possibility, the value of the new item is the unique object in the context set that satisfies the content of the description.[22] Note that the uniqueness requirement is far from absolute. Not only does it allow that in the world there is more than one object that satisfies the content of the description (which absolute quantification would forbid), it even allows that among all the possible values of the discourse items in the state as a whole there are many such objects, also with respect to a single possible world.

Following this recipe, updating the state depicted in figure 2a—the result of updating the sample information state with sentence (3)—with sentence (4), will lead to the state 2c, *via* the intermediary state 2b, which is the result of processing the anaphoric definite 'the man'.

(4) The man is wearing blue suede shoes.

The man that is being talked about has to be Bill, since according to the information of the agent, Bill is the only one wearing blue suede shoes. (But Bill is not the only man, nor is he the only man walking in the park.)

Notice the following. The definite description itself introduces a new discourse item. In the present case, this may seem of little use, since the two discourse items are completely indistinguishable: in each possibility in the information state the two items have the same value. And from here on, they will behave as if they were one and the

19. Accommodation will be left out of consideration in what follows. See Groenendijk et al. 1995a for some discussion.

20. What are called 'pre-conditions' are closely related to presuppositions. For an analysis of presupposition in a dynamic framework see Zeevat 1992; Beaver 1995; Krahmer 1995. For a recent overview of different approaches, see Beaver 1996.

21. For other analyses in a dynamic setting, see Heim 1982; van Eijck 1993; Krahmer 1995.

22. Obviously, this procedure needs further refinement.

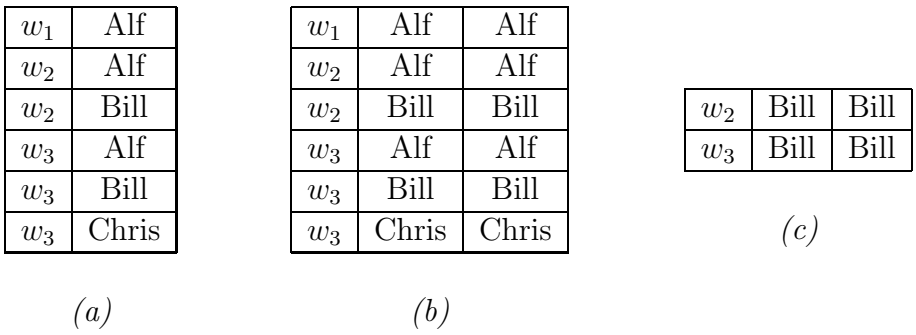

Figure 2: A man walks in the park. *(a)* The man *(b)* wears blue suede shoes. *(c)*

same. We will meet other cases, though, where the introduction of a new item by an (anaphoric) definite description will turn out to be essential.[23]

Notice also that we did not introduce a level of logical (or other) form at which the anaphoric relation is *represented*. To account for anaphoric relations at a level of representation would involve some mechanism of co-indexing. We would have to use the same number, or the same syntactic variable in presenting the contribution of 'a man' and 'the man' to the discourse representation. No mechanism of co-indexing plays a role in the update procedure stated above. The anaphoric definite description picks up its antecedent solely *via* its quantificational force and its descriptive content. Again, in this particular case, one might just as well have used a co-indexing mechanism, linking the definite explicitly with a particular discourse item introduced earlier. However, as we will see shortly, in general the two procedures do make a difference.

3.5 Another man

Not only definite descriptions can be anaphoric, virtually any quantifier can be used in an anaphoric way. The indefinite determiner 'another' is a clear case of a quantifier that can only be interpreted by relating it to context sets. Consider:

(5) A man is walking in the park. Another man is walking in the park, too.

Contextual dependence comes in at several points. First of all, there is the precondition that in every possibility there should be at least one man in the context set of that possibility. If not, the interpretation process comes to a halt. If this precondition is met, the state is extended with a new discourse item, the value of which in a possibility is to be a man from the global domain of discourse, which is not yet a member of the context set of that possibility. How many extensions result in the new state for each old possibility depends on how many such men there are.

Consider again our sample state as it was specified in section 3.2. After an update with the first sentence of (5) it results in the state depicted in figure 3a. A further update with the second sentence of (5) leads to 3c, *via* 3b, which present the effect of processing the anaphoric indefinite 'another man'. Note that world w_1—in which only one man

23. If a state contains two indistinguishable items, this is a good reason for cleansing it by discarding one of the two. Doing so saves space and can make no difference for whatever update is still to follow.

walks in the park—has been eliminated. (Just as w_2 would be eliminated if we repeat the last sentence of (5) once more.) In this case, too, no co-indexing is used to account for the anaphoric link. In fact it is hard to imagine how one could call upon co-indexing as a way to account for this kind of anaphoric relation. (Co-indexing seems particularly unsuited to deal with iterated uses of 'another... (yet) another...'.)

The two discourse items that are present in the information state obtained after processing (5) have a special feature. They are quantitatively distinct: in each possibility they have a different value. But they are qualitatively indistinguishable: for each possibility in which the two items have a particular value, there is another possibility which is the same, except for the fact that the values of the two items are interchanged.[24]

The fact that the items introduced in (5) by the indefinite terms 'a man' and 'another man' are quantitatively different, but qualitatively equal, explains why one cannot refer back to a particular one of the two men involved using a singular anaphoric definite description.[25]

3.6 The one and the other

Of course, it is possible to continue (5) and to refer by anaphoric means to each of the two men separately. One way to do so is as follows:

(6) The one is wearing blue suede shoes, the other is not.

Observe that such anaphoric reference is to neither of the two men in particular. We treat 'the one... the other...' as a polyadic quantifier. Its precondition is that the context set of each possibility consists of two different objects which satisfy the descriptive content

---

24. Continuing the remark made in footnote 23: here one meets another reason for cleansing information states. Since after processing (5), the two discourse items are qualitatively indistinguishable, there is little use in keeping these two separate items. It would do just as well to have a single item instead, the value of which in each possibility is the *set* consisting of the two men in question. This would halve the number of possibilities in state 3c, since the order in which the two have been introduced is irrelevant. Apart from being more economic, such a cleansing operation would make no difference. We abstain from actually performing them, since plural reference is left out of consideration anyway. For extensive discussion of plurality in the context of dynamic semantics see van der Does 1993; van den Berg 1996.

25. Notice the difference between (5) and (i):

(i) A man entered the room. Another man entered the room.

Unlike (5), it is most natural to interpret (i) as a description of two subsequent events. In that case, as participants in two different events, the two men are qualitatively different, which does make it possible to anaphorically refer back to just one of them using a description such as 'the man who entered first' or, simply 'the first' and 'the second'.

Another case in point is:

(ii) Look! A man is walking in the park. Look! Another man is walking in the park, too.

Apparently, both men are located in the visual field of the speech participants, and hence are distinguishable. That is why here, too, a definite description can be used to refer to a particular one of these two men. For example, one could continue (ii) with 'The first one is my brother'. Such a continuation would be out in the case of (5), under the assumption that there is no additional information, visual or otherwise, from outside the discourse that qualitatively distinguishes between the two men.

In the case of (ii) the indefinites are used referentially: for each of the discourse items introduced by them, its value is the same in each possibility, since —by assumption—the object is observationally present. (See Ludlow and Neale 1991; Groenendijk et al. 1996c.)

w_1	Alf
w_2	Alf
w_2	Bill
w_3	Alf
w_3	Bill
w_3	Chris

(a)

w_1	Alf	Bill
w_1	Alf	Chris
w_2	Alf	Bill
w_2	Alf	Chris
w_2	Bill	Alf
w_2	Bill	Chris
w_3	Alf	Bill
w_3	Alf	Chris
w_3	Bill	Alf
w_3	Bill	Chris
w_3	Chris	Alf
w_3	Chris	Bill

(b)

w_2	Alf	Bill
w_2	Bill	Alf
w_3	Alf	Bill
w_3	Alf	Chris
w_3	Bill	Alf
w_3	Bill	Chris
w_3	Chris	Alf
w_3	Chris	Bill

(c)

w_2	Alf	Bill	Alf	Bill
w_2	Alf	Bill	Bill	Alf
w_2	Bill	Alf	Bill	Alf
w_2	Bill	Alf	Alf	Bill
w_3	Alf	Bill	Alf	Bill
w_3	Alf	Bill	Bill	Alf
w_3	Alf	Chris	Alf	Chris
w_3	Alf	Chris	Chris	Alf
w_3	Bill	Alf	Bill	Alf
w_3	Bill	Alf	Alf	Bill
w_3	Bill	Chris	Bill	Chris
w_3	Bill	Chris	Chris	Bill
w_3	Chris	Alf	Chris	Alf
w_3	Chris	Alf	Alf	Chris
w_3	Chris	Bill	Chris	Bill
w_3	Chris	Bill	Bill	Chris

(d)

w_2	Alf	Bill	Bill	Alf
w_2	Bill	Alf	Bill	Alf
w_3	Bill	Alf	Bill	Alf
w_3	Alf	Bill	Bill	Alf
w_3	Bill	Chris	Bill	Chris
w_3	Chris	Bill	Bill	Chris

(e)

Figure 3: A man walks in the park. *(a)* Another man *(b)* walks in the park, too. *(c)* The one ... the other .... *(d)* ... wears blue suede shoes ... does not. *(e)*

of the quantifier, which in this particular case is empty. Thus, the precondition makes use of the only aspect that distinguishes between the two men (in the discourse): that they are quantitatively distinct. If the precondition is met, two new discourse items are added, and for each old possibility, we end up with two new ones: one extension in which in the field of the two new items we find the values of the two old items in the same order, and one in which we find them in the two new fields in the reverse order. (See figure 3d.)

In view of the 'non-specific' nature of the anaphoric reference, it is impossible to co-index one of the elements of the polyadic definite with one of the two preceding indefinites. In the particular case of (5) followed by (6), this may seem of little importance, precisely because the two items introduced by (5) are qualitatively indistinguishable. However, in general this is something to be reckoned with. Consider the following example:

(7) Alf is walking in the park. Bill is walking in the park, too. The one is wearing a hat, the other is not.

When interpreting the last sentence, we can not associate one of the items introduced by the polyadic definite with a specific discourse item, be it either the item introduced by the name 'Alf', or the one associated with the name 'Bill'. To establish such a specific link, we need additional information, i.e., we need to know which of the two actually is wearing a hat. On the other hand, lack of this information does not prevent us from processing this sequence of sentences. If we had to co-index each of the elements of the polyadic quantifier with one particular item in the context, the uninterpretability of this sequence would in fact ensue, which shows that something like the procedure as it was described above, is called for.

Polyadic anaphoric definite descriptions are not the only kind of anaphora that resist linking to specific discourse items. Sometimes also non-polyadic anaphoric definite descriptions behave in this way:

(8) Eva wrote down a number. She wrote down another number. ... She wrote down another number. She subtracted the smallest number from the largest one.

In order to interpret the terms 'the smallest number' and 'the largest number' we need *not* be able to identify particular discourse items as satisfying their descriptive contents. The term 'the largest number' has as its precondition that in each possibility there is among the objects in the context set of that possibility a number which is greater than all others. Analogously for 'the smallest'. (So, both the definite article as such, and the interpretation of 'largest' and 'smallest' involve contextually restricted quantification.) In the example in question, this precondition is easily met.

But, surely, the largest number we find in the one possibility can be the value of one particular item (i.e., can occur in the field in one particular column), whereas the largest number we find in another possibility can be the value of another item (i.e., can occur in the field of another column). It is precisely this feature that blocks an analysis that proceeds by co-indexing the anaphoric definite description with a particular preceding indefinite.

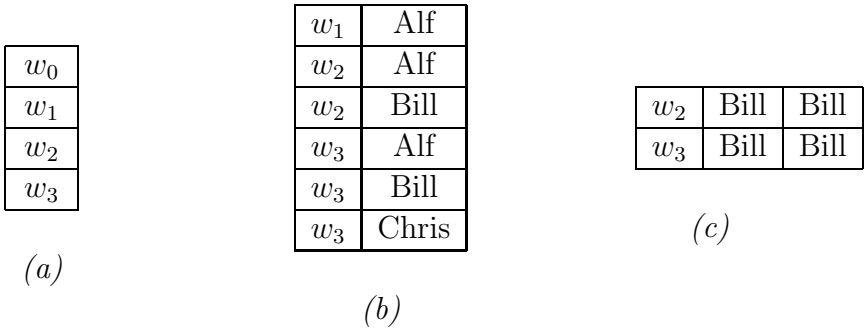

Figure 4: B's states. [Initial state] *(a)* A man walks in the park. *(b)* The man wears blue suede shoes. *(c)*

4 From monologue to dialogue

The examples discussed above all concern (small) monologues of a single speaker, and they were discussed solely from the viewpoint of a hearer. In this section we make some observations concerning the more general case of a discourse with more than one speaker. Again we concentrate on anaphoric relations, which across utterances of different speakers will appear to exhibit special features of interest.[26]

4.1 Paying attention

Before turning to dialogue, it is useful to consider the different roles of speaker and hearer in a monological discourse in some more detail, and introduce some relevant notions. Here, there is one speaker, A, providing information, and one hearer, B, paying attention.

Above we considered the following discourse:

(9) A: A man is walking in the park. The man is wearing blue suede shoes.

and discussed its update effects for a hearer who has the following information. Either no-one is walking in the park, or just Alf, or Alf and Bill, or Alf, Bill and Chris; Bill is wearing blue suede shoes. Assuming B is such a hearer, the update effects on his initial state are as recapitulated in figure 4.

The discourse provides the hearer B with new information. After updating with it, he has the information that Alf and Bill are walking in the park, and that maybe Chris is, too. Furthermore, he has the discourse information that the speaker A must be referring to a particular man, viz., Bill, since he is the only one who is wearing blue suede shoes. In the diagram this corresponds to the fact that the possibility consisting of w_0 is eliminated after an update with the first sentence. The other three initial possibilities then still subsist. After an update with the second sentence, only those possibilities subsist where Bill is the value of the discourse item. This means that the initial possibility

26. Analyses of this type of discourse in a dynamic setting are scarce. See Francez and Berg 1994 for a discussion in the framework of discourse representation theory, and Groenendijk et al. 1996c for some more elaborate discussion along the lines of the present paper.

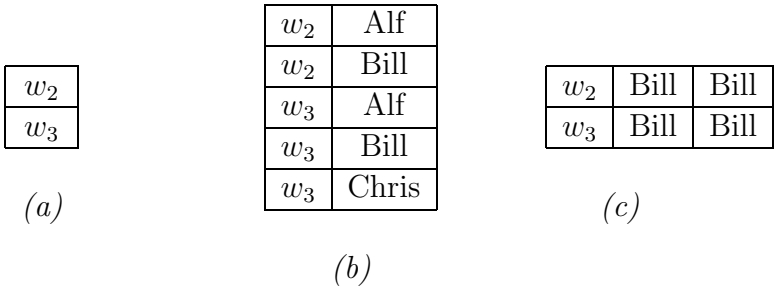

Figure 5: A's states. [Initial state] *(a)* A man walks in the park. *(b)* The man wears blue suede shoes. *(c)*

consisting of world w_1, in which only Alf walks in the park, does not subsist in the final state either.

Observe that the fact that B obtains new information from the discourse, indicates that there is a fundamental difference between speaker and hearer. If a hearer learns something from a discourse this implies that he himself would not have been in the position to utter it sincerely, as we may assume the speaker was. The difference can be explicated in terms of the notion of *support*. For a speaker to utter a sentence correctly it is required that his information state supports it.[27] An information state s supports a sentence ϕ iff every possibility in s subsists after an update of s with ϕ. In other words, for every possibility in s there should be one or more extensions in s updated with ϕ.[28] Clearly, the initial state of the hearer B, as depicted in figure 4, supports neither the first sentence, nor the discourse as a whole, which is why he could obtain new information by updating with it.

A state which does support (9) is the one depicted in figure 5a. It is actually quite like the final state 4c, in which B ended up.[29] It implies that either Alf and Bill, or Alf, Bill and Chris walk; and that Bill wears blue suede shoes. Obviously, if A is in this state, he can sincerely utter (9), for it supports his utterances.

But equally obviously, this state is not the only possible state that supports (9). For example, suppose another speech participant C believes that Alf is not walking in the park, but that Chris is, and that maybe Bill strolls there, too. Suppose furthermore that C believes that Chris is the only man wearing blue suede shoes. Then—even though his state has no possible world in common with B's initial state—his information also supports the utterances in (9). And note that despite this mismatch between the information of C and that of B, the discourse could still proceed without problems, at least up to that point. Of course, if C were to continue the discourse in (9) by saying: 'It is Chris', it would become apparent to B that something is wrong. If B would try to up-

27. This follows Grice's Maxim of Quality.
28. See Groenendijk et al. 1995b; Groenendijk et al. 1996b for more discussion. The notion of support also plays a key role in the definition of dynamic entailment. Roughly, $\phi_1 \ldots \phi_n$ are said to entail ψ iff every state which is updated consecutively with $\phi_1 \ldots \phi_n$ supports ψ.
29. The difference is that in A's initial state 5a, no discourse items occur yet. But after having produced his discourse, A is in the same final state as B: here, exchanging information results in rare close harmony.

date his state 4c with this additional utterance (linking the pronoun with the discourse item present), the result would be that no possibility remains: the *absurd state*. In other words, B's state 4c is inconsistent with the sentence 'It is Chris'.

The notion of (in)consistency is a key notion for a hearer, i.e., for someone who is paying attention. A sentence ϕ is *consistent* with an information state s iff updating s with ϕ does not lead to the absurd state, the state in which no possibility has remained.[30] If the sentence uttered by a speaker is consistent with the information of the hearer, the hearer can update his information with that sentence. If an update with what the speaker has said results in the absurd state, the hearer knows—on the assumption that the speaker utters the sentence sincerely—that his information is incompatible with that of the speaker. Awareness of this fact, will guide him: he will give notice of the observed inconsistency, and a discussion may ensue in order to find out where the difference of opinion lies, and to try and resolve it.

Consistency and support are important semantical notions within dynamic semantics. The first is hearer-oriented, the second is speaker-oriented. The information state of a speaker has to support the sentences he utters in discourse. A hearer will only be willing to update his information state with pieces of discourse which are consistent with his information.

4.2 Exchanging information

Let us now turn our attention towards dialogue, rather than monologue. Consider again the discourse in (9), but suppose that A utters only the second sentence, after its first sentence has been uttered by a different speaker:

(10) D: A man is walking in the park.
A: The man is wearing blue suede shoes.

There is a difference between A's monologue in (9) and the dialogue between D and A in (10). Suppose that before the discourse starts, A's initial state is again the one depicted in figure 5a, which supports 'A man is walking in the park'. After updating with D's utterance, A is in the state 5b. It appears that—although a discourse item is available in A's information state, which seems to license the use of the anaphor 'the man'—A's utterance is infelicitous nonetheless. This is remarkable, in view of the fact that if A were to have uttered the first sentence himself, he could have followed up by uttering the second sentence without problems. After all, as we saw above the monologue is supported by A's information state.

It does not seem too difficult to explain why A's utterance of the second sentence is problematic. If we compare 5b, the state A is in after having updated his initial state 5a with D's utterance, with 5c, the state that results after updating 5b with his own utterance of the second sentence, we see that not all possibilities in 5b subsist in 5c. Those possibilities in which the value of the item in 5b is not Bill, do not subsist in the final state 5c. In other words, 5b, the state of A after updating with D's utterance, does

30. Again, see Groenendijk et al. 1995b; Groenendijk et al. 1996b for some more discussion.

not support his own utterance.[31]

Intuitively, what seems to be at stake is this. It is D who has introduced the discourse item of a man walking in the park. After D's utterance of the first sentence, there are several possible values of the discourse item. D may intend to delimit these possibilities by adding more features to this as yet indefinite man.[32] A, however, does not seem to be licensed to turn D's indefinite man into a more definite one on his own account.[33]

What this observation comes down to, is that there is a difference in acceptability conditions for the use of anaphoric expressions in monological and dialogical situations. It matters for the correct usage of an anaphoric expression who introduced the discourse item(s) it is linked to.[34]

---

31. Observe that also in the monological case, A's second sentence is not supported (in the technical sense) by state 5b, i.e., the state that results after updating A's own initial state with his first sentence. Both the first sentence, and the sequence of the two sentences as a whole, are supported by his initial state, but the intermediate state as it is depicted in figure 5 does not support his second sentence.

32. He may also not intend any such specification, but simply want to draw A's attention to the fact that someone is there; D, or A for that matter, might continue after the first sentence with 'Let's get out of here!'.

33. Of course, if—as in the monological case—A himself has introduced the indefinite man, he is licensed to make him more definite. Continuing footnote 31, that is why there is nothing wrong with A's monologue, even though his intermediate state 5b does not support his second sentence. Continuing footnote 32, it is not unlikely that already before starting his monologue, A *intends* to refer specifically to Bill. The intermediate state in figure 5 does not reflect such intended reference. It reflects who, according to his own information, could be possible referents on the basis of what he has made public himself so far. (By the way, A can never succeed in turning his indefinite man into Chris, i.e., not without loosing support.)

Another example that may point towards the relevance of speaker's intentions is a dialogue-version of the monologue (5), discussed in section 3.5:

(i) D: A man is walking in the park.
A: Another man is walking in the park, too.

If A, as before, is initially in the state depicted in figure 5a, and, hence, in state 5b after having updated with D's utterance, he might seem to be entitled to utter the second sentence in (i). The result of updating A's own state 5b with his utterance would result in the state depicted earlier in figure 3c. Since from the start A has the information that there is more than one man walking in the park, both D's utterance and his own are supported by A's information. Still, A's use of the anaphoric indefinite 'Another man', does not seem to be tremendously felicitous. In this case, the reason is not that A is making D's indefinite man more definite. He is not. He only adds an equally indefinite, qualitatively indistinguishable, but quantitatively distinct man. However, as we observed in section 3.5, when two qualitatively equal man are present in the context sets, one cannot refer back to a particular one of them by using a singular definite description. This means that A's utterance robs D from the possibility to turn 'his' indefinite man into a more definite one.

Another way to look at it is that D may intend to refer to a (more) particular man. To the extent that intentions are 'private', A is not in the position to choose a man who is different from the one possibly meant by D. Except under special circumstances, he has simply no idea whom that might be.

These questions are closely connected to such issues as speaker's reference, and its relation to semantic reference, referential and attributive use, and so on, familiar from the work of Kripke, Donellan, and others. See Dekker 1995 for some discussion in a dynamic setting.

34. This goes against the assumption made in Francez and Berg 1994 that any sequence of sentences that is acceptable as a single speaker discourse, is equally acceptable as a discourse where the different sentences in the sequence are uttered by different speakers.

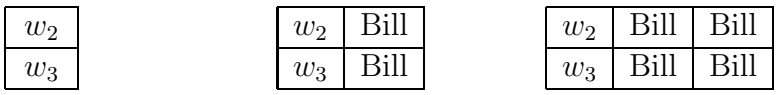

Figure 6: [Initial state] *(a)* Look [pointing at Bill], a man walks in the park. *(b)* Yeah, the man wears blue suede shoes. *(c)*

4.3 Sharing a perspective

We do not claim that the discourse in (10) is unacceptable under all circumstances. Our claim is merely that the acceptability conditions are different in the monological and the dialogical situation. To be sure, (10) *can* be a correct discourse, be it under rather special circumstances. One such case is where D and A take turns in telling (or making up) a story. But this seems to be a mere variant of the monological case, in the sense that under such circumstances D and A are operating as a single agent, rather than as two agents exchanging information.

More interesting is the case where D and A find themselves in a particular kind of observational situation. Suppose that one man is prominently present in the visual fields of both D and A, in such a way that D can be sure that his utterance cannot fail to draw A's attention to this individual. A realizes this, too, and it seems that it is for this reason that he can use the anaphoric definite to refer to this same individual. The following variant of (10) more clearly exposes these features of the utterance situation:

(11) D: Look! A man is walking in the park.
A: Yeah! The man is wearing blue suede shoes.

But if this is the situation, assuming the initial state of A to be the same as depicted in figure 5a, the result of updating it with D's utterance, together with the non-linguistic information provided by D's gestures and possibly other particular features of the situation, results in the intermediate state depicted in figure 6b, rather than the one in figure 5b. In this case, A's own utterance of the second sentence is clearly supported by the state he is in after D's utterance of the first sentence.[35]

4.4 Appreciating the difference

The observations made above may suggest that for the use of an anaphor to be correct it is sufficient if the utterance of the second speaker in a dialogue is supported (in the technical sense) by the information he has after having updated with the utterance of

35. Similar observations can be made concerning the example (i) discussed in footnote 33. Compare this example with:

(ii) D: Look! A man is walking in the park.
A: Yeah! And look! Another man is walking in the park, too.

Here, the utterance situation prevents A from introducing a qualitatively indistinguishable man in the context, and which individual D intends to refer to, is apparently clear to A. (Cf., also footnote 25.)

the first speaker. However, there are several reasons to doubt this.[36]

Consider again the monologue in (9), uttered by A. Suppose B is again in the initial state depicted in 4a. Suppose B reacts to A's utterance as follows:

(12) A: A man is walking in the park. He is wearing blue suede shoes.
B: It is Bill

Clearly, B's utterance is supported by the information state 4c that he is in after having updated with A's utterances.

There may be situations in which this is sufficient and where the exchange is correct. But suppose that B is rather unsure about A's information. Concerning who are walking in the park and who is wearing what kind of shoes A's information might be compatible with his own. But A's information might also be like that of C, who thinks that Chris is the guy wearing blue suede shoes, and who has information about who might and might not be walking in the park which is incompatible with B's own information. If there is such uncertainty about what common knowledge they have, B's use of an anaphor does not seem to be correct just like that. In such a situation, B would rather continue A's utterances as follows:

(13) B: Then it is Bill. (Bill is wearing blue suede shoes.)

The 'Then' in (13) indicates that B draws a conclusion on the basis of his own initial information, updated with what A has been saying. It invites A to check against his own information whether he can share the conclusion or not.[37]

One thing this observation suggests is that if the use of an anaphor by a speaker B in a discourse context created by a speaker A is to be felicitous, it is not only B's own information about the world, and the discourse information linked to that, that counts. The information of the speech participants about the information of each other is equally relevant. Roughly speaking, for B's utterance in (12) to be felicitous, he has to take for granted that there is sufficient consensus about the constitution of the partial object brought under discussion by A to support coordinated co-reference.[38] Lack of sufficient certainty about that does not block B's ability to use anaphora relating to antecedents introduced by A completely, but he has to embed them under an operator like 'Then', which politely invites A to test whether he can agree upon the conclusion B has arrived at concerning the discourse item introduced by A. Leaving the 'Then' out, B would seem to order A bluntly to update with, i.e., to accept what B has figured out for himself about the discourse item introduced by A. The greater the agreement about the object of information A and B assume to share, the more smoothly such unqualified use of the anaphor by B will appear.

---

36. One is the example (i) discussed in footnote 33, where the utterance of the second speaker, containing an anaphoric indefinite 'Another man', was seen to be infelicitous, even though after an update with the sentence of the first speaker, his state supported his own utterance.

37. For other observations and analyses of the dynamic role of such modal expressions see Zeinstra 1990, Vermeulen 1994, chapter 5.

38. The incorporation of such higher order information in the architecture of information states is studied in Groeneveld 1995; Gerbrandy 1996.

4.5 Hearsay

As we described the situation in which B would utter (13) rather than the unqualified sentence in (12), B reckoned with the possibility that A's information was incompatible with his own. However, that is not essential. Also in case B is convinced that A's information is correct, is equally sure about his own information, and has every reason to believe that A's attitude towards his information is no less trustworthy, then the rules of language use still seem to dictate that if B's utterance is supported by his own information updated with what A has said, and not simply on the basis of his own direct information, B should explicitly qualify his utterance as being partially based on what A has said.

Consider the following case. A is visiting B in his apartment, which overlooks a park. It is in the middle of the night. B is preparing another drink in the kitchen. A is looking out of the window, and sees a man in the park in the light of the street lamps. He reports his observation:

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

Based on his long time experience, B knows that always if a man is walking in the park at this time of the night, he is walking his dog. He has no reason whatsoever to distrust A's eyesight. So, on the basis of a simple modus ponens[39] his information state surely supports:

(15) B: He is walking his dog.

But B would not put it like that. He would rather say something like:

(16) B: Then he must be walking his dog.

This invites A to inspect the situation, and respond with something like:

(17) A: Yeah, you're right, he is.

Much in the same way as if B had asked:

(18) B: Is he walking his dog?

The unqualified assertion (15) is only correct, if following A's utterance, B looks out of a window for himself and observes man and dog.[40]

It seems that the rules of discourse are rather strict about this. Independently of how sure we are about our own information, and about the information of the other

---

39. A dynamic one, though.

40. Note that it is not just the potential defeasibility of the B's observational generalization that triggers the 'Then' in his utterance. It is no less needed in the following (rather silly) exchange:

(i) A: The water is boiling.
B: Then it is a 100 degrees Celsius.

Only if B is reading the temperature from a thermometer that is held in the water (imagine that A and B are pupils practicing in a science class) it would be alright for B to say:

(ii) It is a 100 degrees Celsius.

Note, however, that whereas in (16) 'must be' sounds better than 'is', the opposite is true of (i). This, we think, might be related to whether or not defeasibility is taken into account. For an account of defeasible reasoning in a dynamic setting, see Veltman 1996.

speech participants, if we believe ourselves that if ϕ then ψ, we are told that ϕ, and hence can come to the conclusion that ψ, we are not entitled to simply react with ψ, but we have to qualify our utterance of ψ in a way that makes clear that ψ is not supported solely by our own (direct) information, but is a conclusion which is drawn on the basis of our own information together with what the other participant has told us.

Why, one might wonder, are the rules of conversation so cautious about this? The answer, we believe, is that it is a safety measure against the dangers of combining pieces of information from different sources. Person A may be in an information state that supports ϕ, and is consistent with both ψ and 'not ψ'. Person B may be in an information state which is consistent with ϕ, and which supports 'if ϕ then ψ'. A is entitled to assert ϕ. B has no reason not to update with that piece of information. If he does, and treats the new piece of information on a par with his own conditional information that 'if ϕ then ψ', then he arrives in a state which supports ψ. So he would be entitled to utter ψ. Since ψ is consistent with A's information state, there would be no reason for A not to update in turn with this piece of information.

However, had A been aware of the fact that B's justification for saying ψ was that he believes that 'if ϕ then ψ', then he might have been more reluctant to perform the update with B's utterance of ψ. A himself does not believe that 'if ϕ then ψ'. He might actually have good reasons to doubt this. So, had he been aware of the discrepancy in information, he might have started a discussion about it, instead of updating with ψ just like that. The function of a modal qualification as in 'Then ψ' is precisely to make explicit that ψ is a conclusion drawn from the combination of one's own information together with what one has been told. Whereas an unqualified utterance by B of ϕ invites A to update with ϕ if he consistently can, an utterance of 'Then ψ' invites A to test whether his information *supports* ψ. If the test fails, discussion can start about why according to B given that ϕ, ψ has to be the case. In the course of that, A may or may not get convinced by B that ψ.

Another way to put it is that in order to be justified in uttering a non-qualified statement ϕ in a discourse, it should be supported on the basis of own's own direct information. An utterance of 'Then ϕ' is justified if it is supported by one's own direct information, updated with utterances of other participants. The rule seems to be hard and fast, it also has to be obeyed in case one is convinced of the correctness of one's own information and the correctness of the information of the other participants in the discourse. Fortunately, in her profound wisdom language preserves her subjects from their frailties.

5 Summing up

We have described and defended a move from the traditional logical-philosophical concept of meaning as a pictorial relation between language and the world, towards the more progressive view which ties the notion of meaning directly to the process of interpretation of discourse. The bite of it does not lie in metaphors, such as the slogan that

meaning is information change potential, but in providing logical tools to implement and analyze these ideas. However, in the present informal paper these had to remain on the background.

We have shown dynamics at work in an analysis of certain anaphoric relations in terms of contextually restricted quantification. In particular the incorporation of discourse information—next to, and in relation to information about the world—enabled us to actually implement this old idea, where we grant that we showed the drawings of a proto-type, rather than a real machine that could hit the road.

From the observations we made about anaphoric relations across utterances of different speakers, we can draw some conclusions concerning the ways in which the notion of an information state has to be refined in order to be able to make the move from monologue to dialogue. Not surprisingly, the discussion showed that it is essential to extend information states with information of the agent about the information of the other participants in the discourse. Agents have to keep track of which discourse items were introduced by whom. And the conversational right to anaphorically relate to an object of information originating from another participant in the discourse depends on how sure we can be about whether or not we are talking about 'the same thing'. To avoid disinformation to occur too easily, we should keep apart our own direct information, and the information obtained by updating with what others have said. In short, information as it plays a role in interpreting and taking part in dialogue has a much more elaborate structure than exposed in our spreadsheets.

A thin polemic line that ran through the story was an argument against the widespread representationalistic conception of context and interpretation. This conception entered into logical semantics—which is an heir to Frege's anti-mentalism—by way of discourse representation theory. The main motivation in its original presentation in Kamp 1981 was also provided by anaphoric relations, in particular cross-sentential and so-called donkey-anaphora. We hope to have shown that the kind of anaphoric relations discussed above, call for an analysis at the level of information content, rather than by linking formal elements in a language of thought.

Bibliography

Beaver, D.: 1995, *Presupposition and Assertion in Dynamic Semantics*, Ph.D. thesis, University of Edinburgh

Beaver, D.: 1996, Presupposition, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Linguistics*, Elsevier, Amsterdam

Chierchia, G.: 1995, *Dynamics of Meaning*, University of Chicago Press, Chicago

Dekker, P.: 1993a, Existential disclosure, *Linguistics and Philosophy* 16(6), 561–88

Dekker, P.: 1993b, *Transsentential Meditations. Ups and Downs in Dynamic Semantics*, Ph.D. thesis, ILLC/Department of Philosophy, University of Amsterdam, Amsterdam

Dekker, P.: 1995, On context and identity, in J. Groenendijk (ed.), *Ellipsis, Underspecification,Events and More in Dynamic Semantics*, pp 119–46, ILLC, University of Amsterdam, Amsterdam, Dyana-2 deliverable R2.2.C

Francez, N. and Berg, J.: 1994, *A Multi-Agent Extension of DRT*, Technical Report LCL 94-9, Technion - Israel Institute of Technology, Haifa

Gerbrandy, J.: 1996, *Dynamic Epistemic Semantics*, unpublished manuscript

Groenendijk, J. and Stokhof, M.: 1990, Dynamic Montague grammar, in L. Kálmán and L. Pólos (eds.), *Papers from The Second Symposium on Logic and Language*, pp 3–48, Akadémiai Kiadó, Budapest

Groenendijk, J. and Stokhof, M.: 1991, Dynamic predicate logic, *Linguistics and Philosophy* 14(1), 39–100

Groenendijk, J. and Stokhof, M.: 1996, Questions, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Linguistics*, Elsevier, Amsterdam

Groenendijk, J., Stokhof, M., and Veltman, F.: 1995a, Coreference and contextually restricted quantification, in M. Simons and T. Galloway (eds.), *Proceedings from Semantics and Linguistic Theory V*, pp 112–129, CLC Publications, Ithaca, NY

Groenendijk, J., Stokhof, M., and Veltman, F.: 1995b, This might be it, in J. Seligman and D. Westerståhl (eds.), *Proceedings of the 1994 Moraga Workshop*, CSLI, Stanford

Groenendijk, J., Stokhof, M., and Veltman, F.: 1996a, Changez le contexte!, *Langage*

Groenendijk, J., Stokhof, M., and Veltman, F.: 1996b, Coreference and modality, in S. Lappin (ed.), *Handbook of Contemporary Semantic Theory*, pp 179–213, Blackwell, Oxford

Groenendijk, J., Stokhof, M., and Veltman, F.: 1996c, Coreference and modality in multi-speaker discourse, in H. Kamp and B. Partee (eds.), *Proceedings of the Workshop on Context-Dependence in the Analysis of Linguistic Meaning*, IMS, Stuttgart

Groeneveld, W.: 1995, *Logical Investigations into Dynamic Semantics*, Ph.D. thesis, ILLC/Department of Philosophy, University of Amsterdam, Amsterdam

Heim, I.: 1982, *The Semantics of Definite and Indefinite Noun Phrases*, Ph.D. thesis, University of Massachusetts, Amherst, Published in 1989 by Garland, New York.

Heim, I.: 1983, File change semantics and the familiarity theory of definiteness, in R. Bäuerle, C. Schwarze, and A. von Stechow (eds.), *Meaning, Use, and Interpretation of Language*, de Gruyter, Berlin

Hintikka, J.: 1983, *The Game of Language*, Reidel, Dordrecht

Hintikka, J. and Kulas, J.: 1985, *Anaphora and Definite Descriptions: Two Applications of Game-Theoretical Semantics*, Reidel, Dordrecht

Janssen, T. M.: 1996, Compositionality, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Linguistics*, Elsevier, Amsterdam

Kamp, H.: 1981, A theory of truth and semantic representation, in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Formal Methods in the Study of Language*, Mathematical Centre, Amsterdam, Reprinted in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Truth, Interpretation and Information*. Foris, Dordrecht, 1984

Kamp, H.: 1990, Comments on: J. Groenendijk & M. Stokhof 'Dynamic predicate logic', in J. van Benthem and H. Kamp (eds.), *Partial and Dynamic Semantics I*, CCS, Edinburgh, Dyana deliverable R2.1.A

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

Krahmer, E.: 1995, *Discourse and Presupposition*, Ph.D. thesis, Catholic University of

Brabant, Tilburg
Landman, F.: 1986, *Towards a Theory of Information*, Foris, Dordrecht
Ludlow, P. and Neale, S.: 1991, Indefinite descriptions: In defense of Russell, *Linguistics and Philosophy* 14(2), 171–202
Neale, S.: 1993, *Descriptions*, MIT Press, Cambridge, Mass.
Partee, B.: 1996, Montague grammar, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Linguistics*, Elsevier, Amsterdam
van Benthem, J., Muskens, R., and Visser, A.: 1996, Dynamics, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Linguistics*, Elsevier
van den Berg, M. H.: 1996, *Some Aspects of The Internal Structure of Discourse. The Dynamics of Nominal Anaphora*, Ph.D. thesis, ILLC/Department of Computational Linguistics, Amsterdam
van der Does, J.: 1993, *Applied Quantifier Logics*, Ph.D. thesis, ILLC/Department of Philosophy, University of Amsterdam
van Eijck, J.: 1993, The dynamics of description, *Journal of Semantics* 10(2)
Veltman, F.: 1996, Defaults in update semantics, *Journal of Philosophical Logic* 25
Vermeulen, C. F.: 1994, *Explorations of the Dynamic Environment*, Ph.D. thesis, Onderzoeksinstituut voor Taal en Spraak, Utrecht
Westerståhl, D.: 1984, Determiners and context sets, in J. van Benthem and A. ter Meulen (eds.), *Generalized Quantifiers in Natural Language*, pp 45–71, Foris, Dordrecht
Zeevat, H.: 1992, Presupposition and accommodation in update semantics, *Journal of Semantics* 9(4), 379–412
Zeinstra, L.: 1990, Reasoning as discourse, *Master's thesis*, Department of Philosophy, Utrecht University