Coreference and Contextually Restricted Quantification

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

1 Introduction

The aim of this paper is to argue that update semantics is a natural framework for contextually restricted quantification, and to illustrate its use in the analysis of anaphoric definite descriptions and certain other anaphoric terms.

The present discussion remains at an informal level, but takes place against the background of the system of update semantics for the language of modal predicate logic as presented in Groenendijk et al. 1995a; Groenendijk et al. 1995c. The theoretical notions that are used rather casually in the present paper are intended to be in accordance with the formal ones defined in those earlier papers. As for the additions to the system particular to the paper at hand, their formal rendering has to be deferred to another occasion.

On the descriptive level, the paper focusses on (singular) anaphoric definite descriptions. The suggestion made here, is to treat them—together with certain other anaphoric terms—as quantifiers, where quantification is dynamic and contextually restricted.

We share 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. Our contribution to this stock of ideas, is to look upon quantification as being of a dynamic nature—in order to account for binding relations across the syntactic scope of quantifiers—, and, when suitable, restricted to context sets—in order to make sense of the uniqueness preconditions of anaphoric definite descriptions, and the preconditions of other kinds of anaphoric terms.[1]

The point of view that (anaphoric) definite descriptions involve context dependent quantification is not new, of course. We hope to show, though, that

1. The term 'anaphoric' is used here in a liberal way. It is not narrowly restricted to cases where an expression can be linked *via* coindexing with a preceding phrase. It applies to all cases where an expression is used in such a way that its interpretation depends on one or more foregoing phrases.

update semantics allows a natural explication of the way in which the contextually determined domains of quantification come about, an aspect which, combined with its dynamic quantificational mechanism, allows an easy switch between absolute and restricted quantification.

We will also present some arguments against an alternative approach to anaphoric definite descriptions, which accounts for their anaphoric nature by coindexing them with a specific term in the context. We will provide some examples which are intended to show that—at least in some cases—coindexing cannot do the job, whereas contextually restricted quantification can. As it seems to be the case that in those cases where coindexing does work, contextually restricted dynamic quantification can also be used, we venture the hypothesis that the latter is to be preferred as a general mechanism.[2]

However, the empirical field of definites and anaphors 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 take the hypothesis to the test.

2 Context and Information

Update semantics takes a radical stand on the context dependent nature of interpretation. The meaning of a sentence is identified with its context change potential. Contexts are taken to be information states. Hence, meanings are looked upon as update functions on information states. Thus, interpretation creates context. At the same time, the updates constituted by sentences are often partial functions: in order to produce output, they may put constraints on the input. In the dynamic process of interpretation, the making and the making use of the context go hand in hand.

Information states 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. At present, this is the only kind of discourse information we take into account.

Information states are defined as sets of possibilities, where possibilities consist of a possible world, a referent system, and an assignment which links the discourse information represented in the referent system to information about the world.

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.

---

2. We do not claim that *all* anaphoric terms can be treated in this way. In particular, in keeping with previous work, we will assume in what follows that (singular) anaphoric pronouns are treated by means of coindexing, 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.

In an initial information state no discourse has started yet. The possibilities in an initial state only consist of a possible world. The referent system, and hence the assignments, are still empty. They get filled as discourse goes on: discourse items are added to the referent system, and objects are assigned to them as possible referents. Final states are like initial states. Once a discourse is finished, discourse information is of no further use, and can be discarded. Similarly, parts of discourse—even certain parts of a single sentence—can create local discourse information which is erased after the interpretation of that part is completed.

For the purpose of illustration, information states can be depicted as simple matrices, as is shown in the figures below. Each row in the matrix corresponds to a possibility in the information state. The first position in each row is reserved for the possible world in that possibility.[3]

The other columns in the matrix concern discourse information, and are added one by one as the discourse goes on. They can also be deleted again, when the (relevant part of) the discourse is closed off. Each column corresponds to a particular discourse item, introduced by a term in the discourse.[4]

The assignment function in each possibility fills the fields in the corresponding row. Each field contains an object (from the domain of the world in that possibility), which is a possible value of the discourse item in question, given the way things are in the possible world in that possibility, and given the values of the other items in that possibility.

One and the same item may have alternative possible values with respect to one and the same world, and the particular values already assigned to earlier items. This means that one and the same world may figure in two or more different possibilities.[5]

In the figures that follow it is illustrated how information extends as discourse proceeds: In going from one state to the next, new items may be

---

3. 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.

4. We do not take into consideration here that there is also 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.
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 1993 for an analysis of implicit arguments in a dynamic setting.

5. In other papers, we refer to discourse items as 'pegs'. Apart from this terminological variation, what is left out of consideration here is that the referent system also keeps track of which variables of the logical language are in active use, and with which peg they are associated. The matrices also do not show that there is a difference between rows and columns. Information states are *sets* of possibilities. Hence the order of the rows is irrelevant. The possibilities themselves, however, are ordered, in the sense that the order of the columns reflects the order in which the discourse items are introduced. I.e., the discourse items can be identified with their column number. (Indeed, pegs are identified with numbers.)

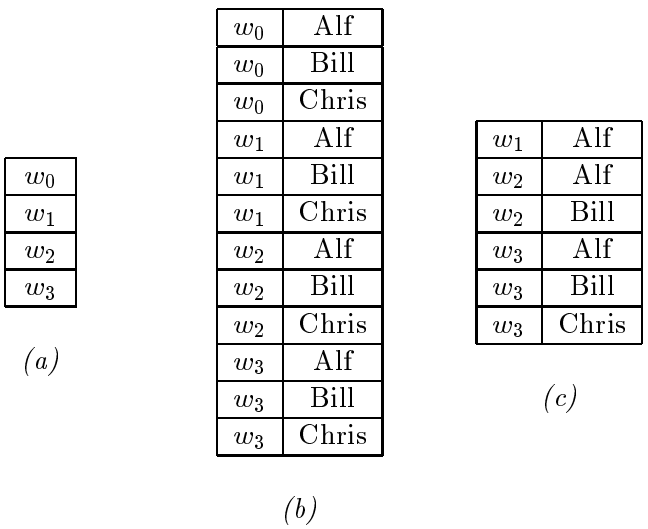

Figure 1: [Initial state] *(a)* A man *(b)* walks in the park. *(c)*

added, where each possibility in the new state is an extension of some possibility already present in the old state. One possibility in the old state may subsist in more than one extension in the new state. But it may also happen that certain possibilities in the old state are eliminated, and do not subsist in the new state. In particular, (all possibilities containing) a world may disappear, which means that we have learned something—albeit not necessarily something true—about the way the world is.

3 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, one has the information that only Bill is wearing blue suede shoes.[6]

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:

(1) A man is walking in the park.

6. It is not that essential to 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. E.g., the way we described 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.

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*.

Indefinites involve existential quantification, one of the dynamic effects of which is the introduction of a new discourse item in the 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.

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

4 Context Sets

As is clear from the way they are depicted, information states naturally come with a contextually restricted domain of discourse. Not only is there in each possibility a global domain of discourse, consisting of all the objects that live in the world of that possibility; but furthermore there is 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.

For example, 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 Westerståhl 1984. Westerståhl stresses the point that a context set is to be distinguished from a universe of discourse. The former, unlike the latter, 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 procedures. In principle there is a choice to be made when one meets a term in the text: 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 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.

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—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.

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. The referent system is extended with a new discourse item, 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.

5 The Man

As for anaphoric definite description, they have as their precondition that within the context set of each possibility, i.e., among the values of the discourse items in that possibility, there is a unique object that satisfies its descriptive content. If this condition can not be fulfilled, the updating proces 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.[7]

Note that the uniqueness requirement is far from absolute. Not only is it not required that in the world there is a unique object that satisfies the content of the description (as absolute quantification would require), even among all the possible values of the discourse items in the state as a whole, there may be 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 (1)—with sentence (2), will lead to the state 2c, *via* the intermediary state 2b, which is the result of processing the anaphoric definite *the man*.

(2) 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 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 same. We will meet other cases, though,

7. Obviously, the procedure as it is described in the text needs further refinement. See the discussion in section 11.

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

(a)

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

(b)

w_2	Bill	Bill
w_3	Bill	Bill

(c)

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

where the introduction of a new item by an (anaphoric) definite description will turn out to be essential.[8]

Notice also that in dealing with the example, no use is made of a coindexing mechanism. 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 have well have used a coindexing mechanism, linking the definite explicitly with a particular discourse item introduced earlier. However, as we will see shortly, in general the two procedures make a difference.

6 Some Donkeys

Heim (Heim 1982, p. 226–9) puts forward the following two examples as *prima facie* problems for a Russellian, i.e., for a quantificational account of (anaphoric) definite descriptions:

(3) If a man beats a donkey, the donkey kicks him

(4) Every boy who likes his mother visits her for Christmas

The difficulty with (3) is how to make sense of the uniqueness that a quantificational approach would require. The problem with (4) is how the definite *his mother* is to bind the pronoun *her*, which is outside its syntactic scope.

Given a dynamic approach to quantification, the second type of example can be dealt with straightforwardly by treating *his mother* as a dynamic quantifier, thus extending its binding force beyond its syntactic scope. Notice by the way that this quantifier is absolute here: for each possible value of the pronominal element which it contains, uniqueness is satisfied in the world, not just relative to a context set.

So, we concentrate on the first type of example. Sentence (3) is a conditional. Processing a conditional involves comparing three states: the input state, the input state hypothetically updated with the antecedent, and the state that results from a further hypothetical update with the consequent.

The update of a conditional sentence as a whole is purely eliminative: the output state will be a subset of the input state. I.e., after having processed

8. 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.

w_0
w_1
w_2
w_3

(a)

w_1	Alf	Heehaw
w_1	Chris	Heehaw
w_2	Alf	Heehaw
w_2	Alf	Eeyore
w_2	Chris	Heehaw
w_3	Alf	Heehaw
w_3	Alf	Eeyore
w_3	Chris	Heehaw

(b)

w_1	Alf	Heehaw	Heehaw	Alf
w_1	Chris	Heehaw	Heehaw	Chris
w_2	Alf	Heehaw	Heehaw	Alf
w_2	Chris	Heehaw	Heehaw	Chris
w_3	Alf	Heehaw	Heehaw	Alf
w_3	Alf	Eeyore	Eeyore	Alf
w_3	Chris	Heehaw	Heehaw	Chris

(c)

w_0
w_1
w_3

(d)

Figure 3: [Initial state] *(a)* [hyp.] a man beats a donkey *(b)* [hyp.] the donkey kicks him. *(c)* If a man beats a donkey, the donkey kicks him. *(d)*

the conditional as a whole, no new discourse items will have been added. New items may be introduced while the procedure is running, but at the end they will have been deleted again. The effect of an update with a conditional is that a possibility is eliminated from the input state unless all its extensions in the state that results after updating with the antecedent, survive a further update with the consequent.[9]

Suppose that an agent has the following information: Either no man beats donkeys, or both Alf and Chris beat Heehaw, and besides, maybe Alf beats Eeyore, too, maybe not. Fortunately, surely no more beating is going on. (Guys wearing blue suede shoes don't do things like that.) Concerning the donkeys, Heehaw is known to be of the kind that kicks back when beaten, about Eeyore information pertaining to his behaviour under such circumstances is lacking.

In case this is the only relevant information, the initial state of the agent reckons with four possible worlds. The state can be depicted as in Figure 3a. (The higher the subscript, the more violence is going on.)

In evaluating (3), we subsequently update this state with the antecedent *a man beats a donkey* and the consequent *the donkey kicks him*. The results are depicted in Figure 3b and 3c, respectively. Testing the initial state with respect

9. It is well known that here are cases of binding across conditionals. But at present, we stick to this oversimplification.

to these two hypothetical states in the way described above, leads to the final state 3d. World w_2—the world in which Alf beats both donkeys, but in which Eeyore is not kicking back—is eliminated. Of the three extensions it had in the state after updating with the antecedent, only two survived a further update with the consequent. Hence, in the final state, those possible worlds will have remained, in which for every man and every donkey such that the man beats the donkey, that donkey kicks that man.

So, using contextually restricted quantification, we meet no problem in interpreting the anaphoric definite description as a quantifier. And note that the kind of uniqueness it requires does not preclude that a man beats more than one donkey. Such possibilities simple survive, provided the vilain is kicked back by every poor beast.[10]

This means that sentence (3), at least when taken as an initial piece of discourse, is equivalent with:

(6) If a man beats a donkey, it kicks him.

Whether or not (5) is fully equivalent with (3), independent of where in the discourse it occurs, is a question which is not easy to answer.[11]

---

10. So, for better or worse, the proposed analysis does not predict strong uniqueness effects. However, what *is* predicted is that if non-uniqueness is *explicitly* communicated, there is a difference. Consider:

(5) Alf is a farmer. He owns a donkey. He owns another donkey. Every farmer who owns a donkey, beats the donkey that he owns.

According to the present analysis, no information state will accept this sequence. The indefinites in the second and third sentences will introduce two donkey-items in the information state, such that in every possibility two different donkeys are assigned to these items. But this means that by the time we meet the anaphoric definite description *the donkey that he owns*, the precondition that there is exactly one donkey among the values of the discourse items is not fulfilled in each possibility. (And neither can it be accommodated.)
So, what is predicted is that informants will have no problems in accepting sentence (3), in case they consider it possible that some men beat more than one donkey, i.e., if uniqueness is not part of the information about the world. However, in case the non-uniqueness has been turned into discourse information, having created possibilities with more than one donkey among the values of the discourse items, informants will judge the sentence to be incorrect in that context.
This is what the analysis predicts. However, for obvious reasons, it seems a prediction that is rather difficult to test by eliciting judgments from informants.

11. From a technical point of view, there are two basic mechanisms available to deal with anaphoric relations within update semantics. The one is coindexing, the other is contextually restricted dynamic quantification, which is what we are discussing here. (Combinations of the two, are also possible. The term *his mother* in example (4) may provide a case in point.)
As far as pronouns are concerned, it is not obvious to us which of the two techniques is to be preferred. If we treat pronouns as quantifiers rather than variables, (3) and (5) can be made fully equivalent. The question that arises is whether they *are* indeed equivalent. For example, if we replace the definite *the donkey that he owns* by *it* in the discourse in the previous note, does that make a difference, or not? Likewise, is there a difference in acceptability (the potential to be accommodated) between the following sequences?

(7) A doctor came in. Another doctor came in. The man said to the woman...
(8) A doctor came in. Another doctor came in. He said to her...

If (15) is judged to be better than (16), this judgment could be used as evidence in favour of treating pronouns by means of coindexing. If one feels little or no difference, that might be evidence in favour of a quantificational approach.

7 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:

(9) 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.

Hence, in our sample state—as it was specified in section 3—,which after an update with the first sentence of (6) results in the state depicted in Figure 4a, a further update with the second sentence of (6) leads *via* 4b, presenting the effect of processing the anaphoric indefinite *another man*, to 4c. Note that world w_1—in which only one man walks in the park—has been eliminated. (Just as w_2 would be eliminated if we repeat the last sentence of (6) once more.)

Note that in this case, too, no coindexing is used to account for the anaphoric link. In fact it is hard to imagine how one could call upon coindexing as a way to account for this kind of anaphoric relation. (Coindexing 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 (6) 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.[12]

The fact that the items introduced in (6) by the indefinite terms *a man* and *another man* are quantitatively different, but qualitatitively equal, explains why one cannot refer back to a particular one of the two men involved using a singular anaphoric definite description.[13]

---

12. Continuing on the remark made in note 8, here one meets another reason for cleansing information states. Since after processing (6), 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 4c, since the order in which the two have been introduced is irrelevant. Apart from being more economic, the effect of such a cleansing operation would make no difference.
We abstain from actually performing such cleansing operations, since the formal system on the background is not yet attuned to plural reference.

13. Notice the difference between (6) and (17).

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

Unlike (6), it is most natural to interpret (17) 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*.

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 4: 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)*

8 The One and the Other

Of course, it is possible to continue (6) and by anaphoric means refer to each of the two men separately. However, such anaphoric reference is to neither of the two men in particular. One way to do so is as follows:[14]

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

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 of the quantifier, which in this particular case is empty. I.e., the precondition makes use of the only aspect that distinguishes between the two men: 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 4d.)

Note that it is impossible to coindex one of the elements of the polyadic definite with one of the two preceding indefinites specifically. In the particular case of (6) followed by (7), this may seem of little importance, precisely because the two items introduced by (6) are quantitatively indistinguishable. However that in general the procedure has to be as it was described above, is obvious from the simpler example:

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

In interpreting the last sentence one can not associate one of the items introduced by the polyadic definite with either the item introduced by the name *Alfred*, or the one associated with the name *Bill*. Unless, that is, we know which of the two actually is wearing a hat. But the lack of this information does not prevent one from being able to process this sequence of sentences. However, if we had to coindex each of the elements of the polyadic quantifier with one particular item in the context, the uninterpretability of this sequence would in

---

Another case in point is:

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

Both men are apparently 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 (18) with *The first one is my brother*. Such a continuation would be out in case of (6), under the assumption that there is no additional information, visual or otherwise, from outside the discourse that qualitatively distinguishes between the two men.

In case of (18) 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 Groenendijk et al. 1995b; Ludlow and Neale 1991.)

14. According to certain stylistic rules for English the *the one... the other* construction is bad, and the *one... and the other* constructions is to be preferred. (In Dutch, by the way, the preference is precisely the other way around.) Unlike the former construction, the latter need not be analyzed as a polyadic quantifier. One can interpret *one* as contextually restricted existential quantification, and *the other* as *the x such that* $x \neq y$, where y is to be the variable introduced by *one*. So, in the analysis of *the other* one has to use both contextually restricted quantification, and coindexing in accounting for its anaphoric impact.

fact ensue.

As a final example of this particular sort, consider yet another familiar donkey:

(14) If a bishop meets another bishop, the one blesses the other

Given the update procedures as they were sketched above, this sentence takes care of itself.

9 Comparing Numbers

Consider the following sequence of sentences:

(15) A man is walking in the park. Another man is walking in the park, too. The man is wearing blue suede shoes.

Obviously, the continuation with the last sentence is infelicitous. Under the analysis proposed here, this is easily accounted for: The uniqueness precondition of the definite is not fulfilled (nor can it be accommodated).

An analysis of anaphoric definite descriptions in terms of coindexing with a preceding term also has to impose a precondition on coindexing to be able to account for the infelicity of (10). If coindexing were free, as seems to be the case in the quantificational approach in van Eijck 1993, nothing would prevent an anaphoric link with just one of the preceding indefinites, rendering (10) ambiguous rather than out. Such problems seem inevitable for any account of anaphoric relations that freely coindexes an anaphor with a *particular* term as its antecedent.[15]

Heim's analysis (see Heim 1982) of anaphoric definites avoids this problem, by formulating a precondition on coindexing. Casted in our terminology (and leaving salience out of the picture), it requires that there is a unique discourse item in the referent system of the information state that satisfies the content of the description (or can be made to do so after accommodation).

Clearly, in case Heim's precondition on the felicitous use of an anaphoric definite description is met, so is ours. But not the other way around. We require there to be a unique *object* in the context set, i.e., among the *values* of all discourse items in each possibility. The objects one finds in different possibilities, may be the values of different items. This freedom is not allowed for in Heim's non-quantificational coindexing approach, which can only link an anaphoric definite description with one particular preceding term.

Above, we already met examples of (polyadic) anaphoric definite descriptions for which it is impossible to make such links with specific earlier introduced discourse items. The following example shares this feature, but it concerns simple, non-polyadic, anaphoric definite descriptions.

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

15. This holds also when quantification is restricted to context sets, as is done in van der Does 1994.

Interpreting the terms *the smallest number* and *the largest number* does *not* require that we 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 other numbers in the context set. Analogously for *the smallest*. (So, both the definite article as such, and the interpretation of *largest* and *smallest* involve contextually restricted quantification.) For 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 (can be in the field in one particular column), whereas the largest number we find in another possibility can be the value of another item (can be in the field of another column). It is precisely this feature that prevents an analysis in terms of coindexing an anaphoric definite description with a particular preceding indefinite.

A Heimian analysis, which amounts to imposing this requirement as a precondition on coindexing an anaphoric definite with a particular preceding indefinite, is hence not able to account for this type of example just like that.

10 The Tallest Man

A similar example, which can be treated in the same way, involves real men instead of mere numbers:

(17) A man is walking in the park. Another man is walking in the park, too. The tallest man is wearing a hat.

It differs from the previous case, in that this example involves accommodation: unlike the relation *larger than* on the domain of numbers, the relation *taller than* on the domain of men is not connected. Accommodation involves the elimination of those possibilities in which the two men in the context set are equally tall.

Suppose that on top of the information the agent had at the beginning of section 3, he furthermore knows that Alf and Bill are equally tall, and that both are taller than Chris. Besides that, he knows that either Alf or Bill is wearing a hat, but not Chris. So, instead of the four worlds depicted in Figure 1a, the initial state of the agent now has to reckon with eight possibilities. Instead of world w_n, we get two worlds w_{na} and w_{nb}, where n reminds us of the number of men walking in the park, and the a or b indicate whether Alf or Bill is wearing a hat.

Figure 5 gives the relevant steps of updating the agents information state with (12). The precondition of *the tallest man* requires that the context set of a possibility contains a man which is taller than the other men in that context set. Accommodating the precondition, no possibility from state 5a in wich the context set consists of Alf and Bill survives in 5b, since Alf and Bill are equally tall. This means that w_{2a} and w_{2b} are altogether eliminated, since there only Alf and Bill walk in the park. And possibilities containing w_{3a} or w_{3b} only survive if Chris is in the context set.

w_{2a}	Alf	Bill
w_{2b}	Alf	Bill
w_{2a}	Bill	Alf
w_{2b}	Bill	Alf
w_{3a}	Alf	Bill
w_{3b}	Alf	Bill
w_{3a}	Alf	Chris
w_{3b}	Alf	Chris
w_{3a}	Bill	Alf
w_{3b}	Bill	Alf
w_{3a}	Bill	Chris
w_{3b}	Bill	Chris
w_{3a}	Chris	Alf
w_{3b}	Chris	Alf
w_{3a}	Chris	Bill
w_{3b}	Chris	Bill

(a)

w_{3a}	Alf	Chris	Alf
w_{3b}	Alf	Chris	Alf
w_{3a}	Bill	Chris	Bill
w_{3b}	Bill	Chris	Bill
w_{3a}	Chris	Alf	Alf
w_{3b}	Chris	Alf	Alf
w_{3a}	Chris	Bill	Bill
w_{3b}	Chris	Bill	Bill

(b)

w_{3a}	Alf	Chris	Alf
w_{3b}	Bill	Chris	Bill
w_{3a}	Chris	Alf	Alf
w_{3b}	Chris	Bill	Bill

(c)

Figure 5: A man walks in the park. Another man walks in the park, too. *(a)* The tallest man *(b)* is wearing a hat *(c)*

The anaphoric definite description itself adds a third item, and in each possibility its field is filled with either Alf or Bill, since they are taller than Chris. Note that the third column is not identical to one of the earlier two. Its value sometimes comes from the one previous column, sometimes from the other. This is precisely why coindexing cannot work in cases like this, and why quantification is called for.

Finally, in updating with the remaining part of the sentence, no further knowledge about the world is obtained, only certain possible assignments are eliminated. However, this additional discourse knowledge may turn out to be of use in case the discourse continues. For example, where the speaker to continue with *He is wearing blue suede shoes, too*, the agent could eliminate the possibilities containing w_{3a}, since only Bill is wearing blue suede shoes. This leaves only one world in the agents information state, world w_{3b}, which means that the agent knows all there is to know about who are walking in the park, and who is wearing what. Provided, of course, that the information he got was correct in the first place.

11 The Man and the Doctor

As we noted, the interpretation procedure for anaphoric definite descriptions outlined in section 5 stands in need of further refinement. Consider the following type of example:

(18) A man came to the doctor. The man said: "...".

Although we may have to reckon with the possibility that *the man* relates to the doctor, by far the most likely interpretation of (13) is that *the man* is anaphorically linked to *a man*. The current treatment does not account for this.

Without accommodation, the sentence would be declared out, since there will be possibilities in which the context set contains two men. Allowing for accommodation, the result is that all possibilities in which the doctor is male are eliminated. For in order to arrive at a unique man in each possibility, as the precondition for anaphoric definite decriptions requires, we would have to infer from the second sentence that the doctor is a woman. Both options certainly are not in accordance with intuitions.

Do examples like (13) then show that the proposed analysis is untenable? We do not think so. But what they do show is that the analysis is in need of further refinement. Information states need to be extended with more structured and detailed representations of discourse information.

In principle there are two ways to assure that the uniqueness condition works in cases such as these. One is to allow only part of the context set to be taken into consideration. The other option is to add further features to the items as such, and to make the uniqueness requirement sensitive to these features. The overall effect will be the same in both cases: it becomes easier to fulfill the uniqueness requirement.

Probably, both strategies are called for. As for the first option, it does seem likely that merely having a list of discourse items is not sufficient. Discourse itself is not just a list of phrases, but has a much more intricate structure. More of this structure should be reflected in the way in which its discourse items are ordered. Consequently, the referent system may consist of different layers of discourse items. One could look upon this as an implementation of part of the notorious notion of salience, discourse items in a higher layer being more salient than ones lower down. Then, the uniqueness precondition need not search through the set of all items, but may be restricted to the items available up to a certain level.

However, it does not seem very likely that the strategy just outlined would work for the type of example (13). It is far from clear that after having processed the first sentence the man is more salient than the doctor. Rather, what seems to be important in this case is that the descriptive contents of the definite *the man* and the indefinite *a man* are much more alike than those of *the doctor* and *the man*: the item corresponding to *a man* is more salient as an object fitting that description, i.e., as a man, than the item corresponding to *the doctor*.

This brings us to the second strategy distinguished above: adding more features of discourse information to the items as such. As things stand, discourse information is not sensitive to the descriptive content by means of which the items are introduced. If this kind of discourse information would be added, the uniqueness precondition could be made sensitive to it, and give the right outcome for cases like (13).

One could then formulate the search procedure in such a way that items that do not fit the contents of the description as well as others may be ignored. The measure of fit can be determined from the values of the items in the in-

formation state as a whole. For example, in case of (13) one will probably find women among the possible values of the item introduced by *the doctor*. And even if all the doctors *are* men (according to the information of the agent), the item introduced by *a man* would still show a better match, as long as not all men are doctors.

Of course, the details need to be spelled out, but it seems not unlikely that if discourse information is refined along these lines, examples such as (13) need not obstruct the kind of quantificational approach proposed here.

A final remark: the strategies outlined above do not necessarily result in deterministic procedures: sometimes equally good solutions for the resolution of an anaphor might result. But then, that seems true to life: even in case of (13) the man may turn out to be the doctor, and not the man.

12 Conclusion

Apart from the empirical differences noted above, Heim's approach and the one proposed in this paper also differ in ontology.

Heim—following Karttunen in this—is a representative of the more general theoretical move to dissociate coreference from reference to the same real object, and to introduce formal objects–discourse referents, file cards, discourse items, pegs—and take coreference to consist in being related to the same such formal object.

The analysis of coreference proposed here, brings coreference back to real reference to objects, within contextually restricted domains. The formal objects, discourse items, still have a role to play, though: their possible values, which are real objects, determine the context sets that quantification can be restricted to.

We hope to have adduced some evidence that the move back from formal coreference (coindexing, really) to real coreference may pay its way, by providing a more adequate and uniform account, not only of anaphoric definite descriptions, but also of other anaphoric elements. However, whether this uniform approach can be maintained throughout remains to be seen.

Acknowledgements

Preparation of this paper was part of the Esprit Basic Research Project Dyana (6852). We would like to thank Maria Aloni, Paul Dekker, Jelle Gerbrandy, Hans Kamp, Tore Langholm, and Craige Roberts for their helpful comments.

Bibliography

Dekker, P.: 1993, Existential disclosure, *Linguistics and Philosophy* 16(6), 561–588

Groenendijk, J., Stokhof, M., and Veltman, F.: 1995a, Coreference and modality, in S. Lappin (ed.), *Handbook of Contemporary Semantic Theory*, Blackwell, Oxford

Groenendijk, J., Stokhof, M., and Veltman, F.: 1995b, *Coreference and Modality in the Context of Multi-Speaker Discourse*, ms

Groenendijk, J., Stokhof, M., and Veltman, F.: 1995c, This might be it, in D. Westerståhl and J. Seligman (eds.), *Language, Logic and Computation: The 1994 Moraga Proceedings*, CSLI, Stanford

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.

Ludlow, P. and Neale, S.: 1991, Indefinite descriptions: In defense of Russell, *Linguistics and Philosophy* 14(2)

Neale, S. (ed.): 1993, *Descriptions*, MIT Press, Cambridge, Mass.

van der Does, J.: 1994, Formalizing E-type anaphora, in P. Dekker and M. Stokhof (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 229–48, ILLC, University of Amsterdam, Amsterdam

van Eijck, J.: 1993, The dynamics of description, *Journal of Semantics* 10(2)

Westerståhl, D.: 1984, Determiners and context sets, in J. v. van Benthem and A. ter Meulen (eds.), *Generalized Quantifiers in Natural Language*, pp 45–71, Foris, Dordrecht