Coreference and Modality

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

1 Static and Dynamic Interpretation

1.1 Towards Dynamic Interpretation

The prevailing view on meaning in logical semantics from its inception at the end of the nineteenth century until the beginning of the eighties has been one which is aptly summarized in the slogan 'meaning equals truth conditions'. This view on meaning is one which can rightly be labeled static: it describes the meaning relation between linguistic expressions and the world as a static relation, one which may itself change through time, but which does not bring about any change itself. For non-sentential expressions (nouns, verbs, modifiers, etc.) the same goes through: in accordance with the principle of compositionality of meaning, their meaning resides in their contribution to the truth-conditions of the sentences in which they occur. In most cases this contribution consists in what they denote (refer to), hence the slogan can be extended to 'meaning equals denotation conditions'.

Of course, although this view on meaning was the prevailing one for almost a century, many of the people who initiated the enterprise of logical semantics, including people like Frege and Wittgenstein, had an open eye for all that it did not catch. However, the logical means which Frege, Wittgenstein, Russell, and the generation that succeeded them, had at their disposal were those of classical mathematical logic and set-theory, and these indeed are not very suited for an analysis of other aspects of meaning than those which the slogan covers. A real change in view then had to await the emergence of other concepts, which in due course became available mainly under the influence of developments in computer science and cognate disciplines such as artificial intelligence. And this is one of the reasons why it took almost a century before any serious and successful challenge of the view that meaning equals truth-conditions from within logical semantics could emerge.

The static view on meaning was, of course, already challenged from the outside, but in most cases such attacks started from premises which are quite alien to the logical semantics enterprise as such, and hence failed to bring about any radical changes.

An important development has been that of speech act theory, originating from the work of Austin, and worked out systematically by Searle and others, which has proposed a radical shift from the proposition with its cognate truth conditions as the principal unit of analysis, to the speech act that is performed with an utterance. Here a move is made from the essentially static relationship between a sentence and the situation it depicts, which underlies the view that meaning equals truth conditions, to a much more dynamically oriented relationship between what a speaker does with an utterance and his environment. This is especially clear from the emphasis that is laid on the performative aspects of speech acts.

This development, however, did not succeed in overthrowing the static logical view, mainly because it turned out not to be a rival, but a companion: the speech act theory of Searle actually presupposes some kind of denotational theory of meaning as one of its components. Nevertheless, speech act theory has been a major influence on work in the logical tradition.

In a similar vein the emergence of the artificial intelligence paradigm only indirectly exercised some influence on the logical tradition. When people working in this area began to think about natural language processing they quite naturally thought of meaning in procedural terms, since, certainly before the development of so-called declarative ('logic') programming languages, the notion of a procedure (or process) was at the heart of that paradigm. This line of thinking, too, may be dubbed dynamic rather than static, since a procedure is essentially something that through its execution brings about a change in the state of a system. However, although this approach has a straightforward appeal, it failed to overthrow the static view, mainly because the way it was worked out failed to address the issues that are central to the logical semantics approach (viz., the analysis of truth and in particular entailment), and also because it lacked the systematic nature that characterizes logical semantics.

The real challenge to the static view on meaning in logical semantics has come from within, from work on recalcitrant problems in logical semantics whose solution required a step beyond the static view on meaning.

Already in the seventies several people had begun to explore a conception of meaning which involved the notion of change. Trying to deal with the many intricacies of context-dependence (such as are involved in presuppositions) Stalnaker suggested that in studying the meaning of an utterance what must be taken into account is the change it brings about in the hearer, more specifically in the information she has at her disposal (see Stalnaker 1974).

Although Stalnaker's conception of meaning has indeed a dynamic, rather than a static flavor, it cannot quite count as a really dynamic notion of meaning after all, for Stalnaker's way of dealing with the dynamic aspect essentially leans on the static conception: he describes the change brought about by the utterance of a sentence in terms of the addition of the proposition the sentence expresses to the set of propositions that constitutes the (assumed) common information of speaker and hearer. But this uses the static notion of a proposition as the basic unit for the analysis of sentence meaning.

In a different setting, that of philosophy of science, Gårdenfors developed dynamic tools (see Gärdenfors 1984) for modeling the structure and change of belief, in particular the process of belief revision.

The real breakthrough, at least within logical semantics, occurred at the beginning of the eighties when, at the same time but independently of each other, Kamp and Heim developed an approach that has become known as 'discourse representation theory' (see Kamp 1981; Heim 1982). Earlier, similar ideas had been put forward within different traditions, such as the work on discourse semantics of Seuren within the framework of semantic syntax (see Seuren 1985), and the work of Hintikka on game-theoretical semantics (see Hintikka 1983). In his original paper, Kamp describes his work explicitly as an attempt to marry the static view on meaning of the logical tradition with its emphasis on truth conditions and logical consequence, with the procedural view emerging from the artificial intelligence paradigm with its appeal of dynamics. Instead of giving it up, both Kamp and Heim stay within the logical tradition in that they want to extend its results, rather than re-do them.

1.2 Dynamic semantics

One particular way of formalizing the idea of dynamic interpretation is the following. It is called 'dynamic semantics' to distinguish it from other approaches, since, as will become clear shortly, it places the dynamics of interpretation in the semantics proper. Unlike other approaches, such as discourse representation theory, which makes essential use of representational structures in the process of dynamic interpretation, dynamic semantics locates the dynamics of interpretation in the very heart of the interpretation process, viz., within the core notions of meaning and entailment.

Very generally, the dynamic view on meaning comes to this: the meaning of a sentence is the change an utterance of it brings about, and the meanings of non-sentential expressions consist in their contributions to this change. This description is general in at least two ways: it does not say what it is that gets changed, and it does not say how such changes are brought about. As in the traditional view, most dynamic approaches start from the underlying assumption that the main function of language is to convey information. Hence, a slightly more concrete formulation can be obtained by replacing in the slogan above 'change' by 'change in information'. But this still leaves a lot undecided: what is this information about, and whose information is it? Here, the empirical domain that one is concerned with gets to play a role. For example, when one analyzes anaphoric relations between noun phrases and pronominal anaphors, the relevant information is that of the hearer about individuals that have been introduced in the domain and about the binding and scope relations that obtain between them. When analyzing temporal relations in discourse, information concerns events, points in time, and such relations between them as precedence, overlap, and so on. In other cases, for example when describing information exchanges such as question–answer dialogues, the information that is relevant is about the world, and one has to keep track of both the information of the questioner and that of the addressee. When analyzing the way presuppositions function in a discourse, another aspect is introduced: the information which the speech participants have about each other's information.

Leaving these distinctions and refinements aside, and restricting ourselves to sentences, the dynamic view can be paraphrased as follows: 'meaning is information change potential'. Per contrast, the static view can be characterized as: 'meaning is truth conditional content'.

In line with this difference, it must be observed that in a static semantics the basic notion that occurs in the definition of interpretation is that of information *content*, whereas in a dynamic system it is the notion of information *change* that is defined recursively. As is to be expected, different views on meaning lead to different views on entailment. In a static system entailment is meaning inclusion. In a dynamic system there are several options. One that is rather natural is the following: a premise entails a conclusions iff updating any information state with the premise leads to an information state in which the conclusion has to be accepted.

The remainder of this paper is devoted to an analysis of a specific problem area, which is not only of interest descriptively, but which also presents an interesting theoretical challenge.

The descriptive area is that of the interaction between indefinites, pronouns, and epistemic modalities, a subject renowned for the many puzzles it creates, including questions concerning identity of individuals, specificity of reference, and rigidity of names. Obviously, not all of these long-standing problems can be studied in depth within the span of a single paper. The aim is merely to show that the dynamic perspective suggests interesting new solutions to some of them.

The paper provides a dynamic semantics for a language of first order modal predicate logic. This system is meant to combine the dynamic semantics for predicate logic developed in Groenendijk and Stokhof 1991 with the update semantics for modal expressions of Veltman to appear. This combination is not a straightforward fusion of two distinct systems, but poses some interesting technical problems. Various people have studied this issue (see van Eijck and Cepparello to appear; Dekker 1992), and the present paper builds on their work. It tries to solve the problems in a different way, by slightly adapting the original definition of existential quantification in dynamic predicate logic, and making use of the notion of a referent system, originally developed in Vermeulen to appear b.

Natural language is not the primary target of the analyses provided below. However, it is a main source of inspiration, and the paper claims that the dynamic approach which is exemplified here using a logical language, can be applied fruitfully to natural language, too. The long term aim is to come up with a logical system which may function as a tool in the analysis of natural language meaning in much the same way as Montague's IL. The present paper is meant as a step towards that goal.

2 Information

In dynamic semantics the meaning of a sentence is equated with its potential to change information states. An implementation of this idea requires, among other things, a specification of the nature of information states. One general conception of an information state is that of a set of possibilities, consisting of the alternatives which are open according to the information. The nature of the possibilities that make up information states depends on what the information is about.

2.1 Two Kinds of Information

Not all discourse serves the same purpose. Here, the focus on one such purpose: that of information exchange. Within this (limited) perspective, two kinds of information need to be distinguished.

First, there is factual information, i.e., *information about the world.* In the end, that is what counts: to get as good an answer as possible to the question what the world is like is the prime purpose of this type of discourse.

There are many ways in which information about the world can be gathered: through perception, reasoning, recollection. One particular way is by the use of language: linguistic communication. And this is what is at stake here: the interpretation of informative language use. This type of discourse is primarily focussed on answering questions about the world. But the interpretation process brings along its own questions.

When one is engaged in a linguistic information exchange, one also has to store *discourse information.* For example, there are questions about anaphoric relations that need to be resolved. This requires a mechanism to keep track of the objects talked about and the information gathered about them; a model of the information of other speech participants has to be maintained; and so on.

In the present paper the focus will be on discourse information of the first kind. Discourse information of this type looks more like a book-keeping device, than like real information. Yet, it is a kind of information which is essential for the interpretation of discourse, and since the latter is an important source of information about the world, discourse information, indirectly, also provides information about the world.

Information About the World

Information about the world is represented as a set of possible worlds, those worlds that given the information that is available still might be the real one. Worlds are identified with complete first order models. Such models consist of a set of objects, the domain of discourse, and an interpretation function. Relative to the domain of discourse, the interpretation function assigns a denotation to the non-logical vocabulary of the language, individual constants and predicates.

In this paper it is assumed that language users know which objects constitute the domain of discourse (although they may not know their names). Consequently, all possible worlds share one domain.[1] Hence, a pos-

1. In due course, this is an assumption one would like to drop. For normally, one is only

sible world can be identified with the interpretation function of a first order model.

Since they are identified with (interpretation functions of) complete first order models, worlds are 'total' objects. Information of language users about the world is characteristically partial. Partiality of information about the world is accounted for by representing it as a set of alternative possibilities. Extending this kind of information amounts to eliminating worlds which were still considered possible.

Even taking into account the restriction to a first order language, this picture of information is very simple, and in many ways not 'realistic'. An obvious alternative is to look upon information as a partial model of the world which is gradually extended as information grows. Or one might combine this more constructive approach with the eliminative approach as it was sketched above. Such more constructive approaches, however, are technically more complicated. For the purpose of the present paper, it suffices to explore the simplest, the eliminative approach.

Discourse Information

As was said above, discourse information keeps track of what has been talked about. In the logical language at hand, it is the use of an existential quantifier that introduces a new item of conversation, a new *peg*. Pegs are formal objects. One can think of them as addresses in memory, for example. But it does not really matter what pegs are. The only thing that counts is that they can be kept apart, and that there are enough of them, no matter how many things are introduced in a discourse. In what follows, natural numbers will be used as pegs. Pegs are introduced one by one in consecutive order, starting from 0.

Variables are the anaphoric expressions of the logical language. To enable the resolution of anaphoric relations, discourse information also keeps track of the variables which are in use, and the pegs with which they are associated. The use of a quantifier $\exists x$ adds the variable x to the variables that are in active use; it introduces the next peg, and associates the variable x with that peg. This is how discourse information grows: extending discourse

partially informed about what there is. No deep technical issues are involved, the reason for the choice made here is convenience. As a matter of fact, the system outlined in the present paper deals with one particular way in which information is partial. There are many others, which are equally interesting, but not all of the same nature. Some of these will be pointed out along the way. However, it will not do to try to deal with them all at once in the scope of a single paper.

information is adding variables and pegs, and adjusting the association between them.

Linking the Two Kinds of Information

Gathering discourse information is not an aim in itself, it is to serve the purpose of gathering information about the world. To achieve this, discourse information is connected to information about the world. The two kinds of information are linked *via* possible assignments of objects from the domain of discourse to the pegs (and hence, indirectly, to the variables associated with these pegs). In general, not every assignment of an object to a peg is possible — both the discourse and the information that is available may provide restrictions —, but usually, more than one is.

Getting better informed on this score is eliminating possible assignments. Suppose a certain assignment is the only one left with respect to some world which is still considered possible. In that case elimination of the assignment brings along the elimination of the world. This is how discourse information may provide information about the world.

2.2 Information States

Referent Systems

In the possibilities that make up an information state, the discourse information is encoded in a *referent system*,[2] which tells which variables are in use, and with which pegs they are associated:

Definition 2.1 A *referent system* is a function r, which has as its domain a finite set of variables v, and as its range a number of pegs.

If the number of pegs in a referent system is n, then the numbers $m < n$ are its pegs.

The use of a quantifier $\exists x$ adds the variable x to the variables that are in use, it introduces the next peg, and associates the variable x with that peg. The corresponding update of a referent system is defined as follows:

Definition 2.2 Let r be a referent system with domain v and range n. $r[x/n]$ is the referent system r' which is like r, except that its domain is $v \cup \{x\}$, its range is $n+1$, and $r'(x) = n$.

2. The use of referent systems was inspired by the work of Kees Vermeulen. See Vermeulen to appear b, Vermeulen 1994a, chapter 3.

Note that it is not excluded that x is already present in v. This situation occurs if the quantifier $\exists x$ has been used before. In that case, even though the variable x was already in use, it will be associated with a new peg. The peg that x was connected with before remains, but is no longer associated with a variable. This means that a referent system r is an injection.

The main reason to allow for the possibility to re-use a quantifier, is that this is usual logical practice. But a case can be made that in natural language things work in a similar way. A noun phrase such as 'a man' introduces a new peg associated with that noun phrase. A subsequent anaphoric pronoun 'he' would be linked to that same peg. If later on in the discourse the noun phrase 'a man' is used again, it should introduce a new peg and associate it with this occurrence of the noun phrase, and a subsequent anaphoric pronoun 'he' would naturally be linked to this new peg. One could still refer back to the first man, but not by using a pronoun, but rather by means of a definite description, such as 'the man I talked about earlier'.

Associating a variable with a new peg is the prototypical way in which the discourse information encoded in a referent system is extended:

Definition 2.3 Let r and r' be two referent systems with domain v and v', and range n and n', respectively.
r' is an *extension* of r, $r \leq r'$, iff $v \subseteq v'$; $n \leq n'$; if $x \in v$ then $r(x) = r'(x)$ or $n \leq r'(x)$; if $x \notin v$ and $x \in v'$ then $n \leq r'(x)$.

A referent system r' is an extension of r iff (i) the variables which were in use in r are still in use in r', but new variables may have been added to r'; (ii) r' has as least as many pegs as r; (iii) the variables that were in use already in r either remain associated with the same pegs in r', or are associated with new pegs, just as (iv) the variables in r' which were not already in use in r are associated with new pegs.

Note that a referent system $r[x/n]$, as defined above, is always a *real* extension of r.

Possibilities

Above a distinction was made between discourse information, information about the world, and a link between the two. These three ingredients are present in the *possibilities*, which in turn make up information states.

Definition 2.4 Let D, the *domain of discourse*, and W, the set of *possible worlds*, be two disjoint non-empty sets.
The *possibilities* based on D and W is the set I of triples $\langle r, g, w \rangle$, where r is a referent system; g is a function from the range of r into D; $w \in W$.

The function g assigns an object from the domain of discourse to each peg in the referent system. The composition of g and r indirectly assigns values to the variables that are in use: $g(r(x)) \in D$.

The possibilities contain all that is needed for the interpretation of the basic expressions of the language: individual constants, variables, and n-place predicates.

Definition 2.5 Let α be a basic expression, $i = \langle r, g, w \rangle \in I$, with v the domain of r, and I based upon W and D.
The *the denotation of* α in i, $i(\alpha)$, is defined as:

i. If α is an individual constant, then $i(\alpha) = w(\alpha) \in D$.
ii. If α is an n-place predicate, then $i(\alpha) = w(\alpha) \subseteq D^n$.
iii. If α is a variable such that $\alpha \in v$, then $i(\alpha) = g(r(\alpha)) \in D$, else $i(\alpha)$ is not defined.

The first two clauses exploit the identification of possible worlds with interpretation functions of first order models. If $i(c) = w(c) = d$, this means that in world w in the possibility i the denotation of the name c is the object d. Similarly for predicates.

The value of a variable is determined by the referent system and the assignment. It is the object assigned by g to the peg that is associated with x by the referent system r. Recall that variables are anaphors, hence they need antecedents: a variable will only be assigned a value if it has already been introduced in the domain of the referent system.

Information States

Information states are subsets of the set of possibilities:

Definition 2.6 Let I be the set of possibilities based on D and W.
The set of *information states* based on I is the set S such that $s \in S$ iff $s \subseteq I$, and $\forall i, i' \in s{:}\, i$ and i' have the same referent system.

Variables and pegs are introduced globally with respect to information states. That is why an information state has a unique referent system. Instead of putting a copy of this single referent system in each possibility, it could also be introduced as a separate component. However, the present set-up makes the definitions run more smoothly, and for the language of modal predicate logic there is no difference.[3]

An information state encodes information about the possible denotations of the expressions of the language. For example, the question who c is, what the denotation of the name c is, is settled in an information state if in all worlds in the information state the denotation of c is the same. And, similarly, the question which objects have the property P is answered if in all worlds the denotation of P is the same. Note that in order to have the information that c has the property P, the questions who c is and which objects have P, need not be settled. It suffices that in each world in each possibility the denotation of c in that world is in the denotation of P in that world.

An information state also encodes information about the possible values of variables. This information, too, is relative to possible worlds. Consider the existentially quantified formula $\exists x Px$. This conveys the information that there is an object which has the property P. Updating an information state with this formula results in a state s in which the following holds: in every possibility $i = \langle r, g, w \rangle$ in s the assignment g will assign to the peg associated with x by r an object which in w has the property P. If there is more than one object in w with the property P, then there will be several alternative possibilities i' with the same world w, assigning different objects in the denotation of P in w to the peg associated with x.

Thus, the typical situation is one in which the same world appears in several possibilities, which differ in the assignments of objects to the pegs. With respect to the same world there may be different possible assignments of objects to the pegs.

2.3 Information Growth

In dynamic semantics information states are used to define the information change potential of expressions. The change brought about by (the utterance of) a sentence defines a relation between information states. Among the various relations between such states, the relation of *extension* (or strengthening) is of primary importance.

---

3. It could turn out to be the case that for a proper account of phenomena such as ambiguity, different alternatives regarding discourse information need to be distinguished, which would surface as different referent systems in distinct possibilities.

Assignment

One way in which information states can be extended is by adding variables and pegs to the referent system, while assigning some object to them:

Definition 2.7 Let $i = \langle r, g, w\rangle \in I$; n the range of r; $d \in D$, $s \in S$.
 i. $i[x/d] = \langle r[x/n], g[n/d], w\rangle$.
 ii. $s[x/d] = \{i[x/d] \mid i \in s\}$.

According to the second clause, assigning some object d in the domain of discourse to a variable x in an information state s is a pointwise operation on the possibilities i in s. Given definition 2.2, the first clause boils down to this: the next peg is added to the referent system r of i, the variable x is associated with this peg, and the object d is assigned to it.

It will appear that this assignment procedure plays an important role in the interpretation of existential quantification.

Extension

Information can grow in two ways: by adding discourse information, and by eliminating possibilities. Both are captured in the following definition:

Definition 2.8 Let $i, i' \in I, i = \langle r, g, w\rangle$ and $i' = \langle r', g', w'\rangle$, and $s, s' \in S$.
 i. i' is an *extension* of i, $i \leq i'$ iff $r \leq r'$, $g \subseteq g'$, and $w = w'$.
 ii. s' is an *extension of* s, $s \leq s'$ iff $\forall i' \in s' \colon \exists i \in s \colon i \leq i'$.

An information state s' is an extension of state s if every possibility in s' is an extension of some possibility in s. This means that in the new state some of the possibilities of the original state may have disappeared. These are properly eliminated. Other possibilities, or one or more extensions of them, may re-occur.[4]

A possibility i' is an extension of a possibility i if i' differs from i at most in that in i' variables have been added and associated with newly introduced pegs, that have been assigned some object.

4. It is worth noting that the definition of the extension relation between information states is largely independent of the particular contents of what constitutes the set of possibilities. For example, whether worlds are taken to be total objects, as is the case here, or partial ones, does not make a difference. Also, incorporating 'higher-order' information as such does not necessitate a change in the definition of extension. Generally speaking, what counts is that possibilities are constructed from sets of objects which are (partially) ordered.

A simple example. Suppose an information state s is updated with the sentence $\exists xPx$. Possibilities in s in which no object has the property P will be eliminated. The referent system of the remaining possibilities will be extended with a new peg, which is associated with x. And for each old possibility i in s, there will be just as many extensions $i[x/d]$ in the new state s', as there are objects d which in the possible world of i have the property P. So, it may very well happen that even though some possibilities are eliminated, the number of possibilities in s' is larger than in s. Still, each possibility in s', will be an extension of some possibility in s. Actually, every i' in s' will be a real extension of some i in s, hence s' will be a real extension of s.

If the resulting state s' is subsequently updated with the atomic formula Qx, then all possibilities in s' will be eliminated in which the object assigned to the peg associated with x does not have the property Q. So, in this case, the resulting state s'' is just a subset of s': there is only elimination of possibilities, no extension of them.

The extension relation is a partial order. There is a unique minimal information state, the *state of ignorance*, in which all worlds are still possible and no discourse information is available yet. This state, $= \{\langle \emptyset, \emptyset, w\rangle \mid w \in W\}$, is referred as $\mathbf{0}$. Subsets of the state of ignorance are called *initial states*. In such states there may be some information about the world (some possible worlds are eliminated), but there is no discourse information yet. The maximal element in the extension ordering, $\mathbf{1} = \emptyset$, is called the *absurd state*. It is the state in which no possibility is left. Less maximal, but more fortunate, are states of *total information*, consisting of just one possibility.

Subsistence

Some auxiliary notions, which will prove useful later on, are the following:

Definition 2.9 Let $s, s' \in S, s \leq s', i \in s, i' \in s'$.
  i. i' is *a descendant of i in s'* iff $i \leq i'$.
  ii. *i subsists in s'* iff i has one or more descendants in s'.
  iii. *s subsists in s'* iff all $i \in s$ subsist in s'.

It follows from the definition that if s subsists in s', then s' is an extension of s'. This means that every possibility in s' is an extension of some possibility in s. But if s subsists in s', it also holds that no possibility in s is eliminated. The state s' may contain more information than s, but only in the sense that variables and pegs may have been added and have been assigned some object. That is to say, whatever new information s' contains, is discourse information, not information about the world. If two states have the same referent system, then the one can only subsist in the other if they are identical.

By way of illustration, consider again the update of s with $\exists x Px$. The original state s subsists in the resulting state s' if there are no possibilities in s in which no object has the property P. In that case every i in s subsists in s', and there will be as many descendants $i[x/d]$ of i in s' as there are objects d in $i(P)$.

3 Updating Information States

Information states being defined, they can be put to use in providing a dynamic interpretation for the language of modal predicate logic.

A formula ϕ of this language is interpreted as a (partial) function, $[\phi]$, from information states to information states. Postfix notation is used: $s[\phi]$ is the result of updating s with ϕ, $s[\phi][\psi]$ is the result of first updating s with ϕ, and next updating $s[\phi]$ with ψ. Whether s can be updated with ϕ may depend on the fulfillment of certain constraints. If a state s does not meet them, then $s[\phi]$ does not exist, and the interpretation process comes to a halt.

Definition 3.1 Let $s \in S$ be an information state, and ϕ a formula of the language. The *update of s with ϕ* is recursively defined as follows:

i. $s[Rt_1 \ldots t_n] = \{i \in s \mid \langle i(t_1), \ldots, i(t_n)\rangle \in i(R)\}$.
ii. $s[t_1 = t_2] = \{i \in s \mid i(t_1) = i(t_2)\}$.
iii. $s[\neg\phi] = \{i \in s \mid i \text{ does not subsist in } s[\phi]\}$.
iv. $s[\phi \wedge \psi] = s[\phi][\psi]$.
v. $s[\exists x\phi] = \cup_{d \in D}(s[x/d][\phi])$.
vi. $s[\Diamond\phi] = \{i \in s \mid s[\phi] \neq \emptyset\}$.

The update of an information state with an atomic formula eliminates those possibilities in which the objects denoted by the arguments do not stand in the relation expressed by the predicate. The same holds for identity statements: those possibilities are eliminated in which the two terms do not denote the same object.

The update expressed by an atomic formula may be partial. If one of the argument terms of the formula is a variable that is not present in the referent system of the information state to which the update is to be applied, then its denotation is not defined, and hence the update does not exist. This source of undefinedness percolates up to all the other update clauses. If somewhere in the interpretation process a variable occurs that at that point has not been introduced, then the whole process comes to a halt.[5]

In calculating the effect of updating a state s with $\neg\phi$, s is updated hypothetically with ϕ. Those possibilities that subsist after this hypothetical update are eliminated from the original state s.

Updating a state with a conjunction is a sequential operation: the state is updated with the first conjunct, and next the result is updated with the second conjunct. The update expressed by a conjunction is the composition of the updates associated with its conjuncts.

If a state s is updated with $\exists x\phi$, its referent system is extended with a new peg, and the variable x is associated with that peg. An object d is selected from the domain and assigned to the newly introduced peg. Then the state $s[x/d]$ is updated with ϕ. This procedure is repeated for every object d. The results are collected, and together make up the state $s[\exists x\phi]$.

The operator $\Diamond$ corresponds to the epistemic modality *might*. Updating a state s with $\Diamond\phi$ amounts to testing whether s can be consistently updated with ϕ. If the test succeeds, the resulting state is s again. If the test fails because updating s with ϕ results in the absurd state, then $s[\Diamond\phi]$ is the absurd state.

The semantics just presented defines the interpretation of the formulae of the language in terms of their information change potential. Actually, they change information states in a particular way:

Fact 3.1 For every formula ϕ and information state s: $s \leq s[\phi]$.

In view of this observation the semantics defined above can properly be called an *update semantics*. The interpretation process always leads to an information state that is an extension of the original state.

Other logical constants can be introduced in the usual way. Calcula-

5. This source of undefinedness is a kind of presupposition failure. The presupposition is of a particular nature: the condition that a variable be introduced, can not be expressed in the object language. It is a 'meta'-presupposition concerning discourse information. Also, it cannot be accommodated.

tion of the definitions results in:

Fact 3.2

i. $s[\phi \to \psi] = \{i \in s \mid$ if i subsists in $s[\phi]$, then all descendants of i in $s[\phi]$ subsist in $s[\phi][\psi]\}$.
ii. $s[\phi \vee \psi] = \{i \in s \mid i$ subsists in $s[\phi]$ or i subsists in $s[\neg\phi][\psi]\}$.
iii. $s[\forall x\phi] = \{i \in s \mid$ for all $d \in D$: i subsists in $s[x/d][\phi]\}$.
iv. $s[\Box\phi] = \{i \in s \mid s$ subsist in $s[\phi]\}$.

It is not possible to make a different choice of basic and defined constants which leads to the same overall results. This can be seen as follows. From the definitions of negation and the existential quantifier it follows that an existential quantifier which occurs inside the scope of a negation cannot bind variables outside its scope.

Consider $\neg\exists xPx$. The negation eliminates all possibilities in which the denotation of P is non-empty. The change in the referent system of an information state s that takes place inside the hypothetical update of s with $\exists xPx$ which is performed in calculating the update of s with $\neg\exists xPx$, is not inherited by the resulting state $s[\neg\exists xPx]$. Hence, the existential quantifier no longer acts dynamically. This means that in general $\exists x\phi$ cannot be defined as $\neg\forall x\neg\phi$. For similar reasons conjunction cannot be defined in terms of negation and disjunction, or negation and implication.[6]

3.1 Consistency, Support, and Entailment

Truth and falsity concern the relation between language and the world. In dynamic semantics it is information about the world rather than the world itself that language is related to. Hence, the notions of truth and falsity cannot be expected to occupy the same central position as they do in standard semantics. More suited to the information oriented approach are the notions of *consistency* and *support*.

Two very simple observations concerning information exchange illustrate this. No hearer will be prepared to update his information state with a sentence if the result would be the absurd state. And a speaker can only assert a sentence correctly if it does not constitute a 'real' update in her information state.

6. In some contexts, negation in natural language does not block anaphoric reference. See Groenendijk and Stokhof 1991, pp. 89–92, Groenendijk and Stokhof 1990a, pp. 37–47, Dekker 1993b, chapter 2 for discussion.

Definition 3.2 Let s be an information state.

i. ϕ is *consistent* with s iff $s[\phi]$ exists and $s[\phi] \neq \emptyset$.

ii. ϕ is *supported* by s iff $s[\phi]$ exists and s subsists in $s[\phi]$.

If ϕ is *consistent* with s, this is often expressed by saying that s *allows* ϕ. The notion of support is defined in terms of subsistence, and not, as one might perhaps expect, in terms of identity. If $s = s[\phi]$ then s supports ϕ, but the converse does not hold.

Consider a (non-absurd) information state s such that there is no possibility in s containing a world in which the denotation of P is empty. Intuitively, such an information state should count as one which supports the sentence $\exists x Px$. Nevertheless, as was indicated already above, in such a state s it will never be the case that $s = s[\exists x Px]$, it will always hold that $s[\exists x Px]$ is a *real* extension of s. However, the information added to such a state s is purely discourse information. In updating such a sate s with $\exists x Px$ no possibility, no possible world or possible assignment of objects to the pegs already present in s is eliminated. This is precisely what is captured in the notion of subsistence, and hence in the notion of support.

Consistency and Coherence

The observations about information exchange made above can be generalized as follows. A sentence is unacceptable if there is not at least *some* state with which is consistent. And if a sentence is not supported by *any* non-absurd state, which means that no speaker could ever sincerely utter it, then that sentence is judged unacceptable, too.

Definition 3.3

i. ϕ is *consistent* iff there is some information state with which ϕ is consistent.

ii. ϕ *is coherent* iff there is some non-absurd state by which ϕ is supported.

Note that coherence implies consistency. Concerning the acceptability of a single sentence, it would suffice to require coherence. Still, it is important to distinguish between these two notions if not just the acceptability of single sentences is at stake, but the acceptability of a discourse, which may consist of a sequence of sentences possibly uttered by different speakers in different information states.

The acceptability of such a discourse minimally requires that one by one the sentences which make it up are coherent. But it would be wrong to require that the discourse as a whole can be supported by a single information state.

For example, one speaker might start a discourse uttering the sentence 'It might be raining outside'. And a hearer may be able to happily confirm this. Another speaker, or even the same one after having opened the blinds, can continue the discourse with 'It isn't raining'. And the same hearer could easily be able to consistently update his state with this information. So, the discourse as a whole is consistent. And each of its two sentences taken separately is coherent. Yet, without an intermediate change in information state, as can be caused by opening the blinds and looking outside, no single speaker can coherently utter the discourse as a whole.

Entailment

The properties of consistency and coherence present criteria for testing the adequacy of a proposed semantics. For the same purpose, the notion of entailment is important, too. Entailment is not defined in the usual way in terms of truth, but in terms of sequential update and support:[7]

Definition 3.4 $\phi_1, \ldots, \phi_n \models \psi$ iff for all information states s such that $s[\phi_1]\ldots[\phi_n][\psi]$ exists, it holds that $s[\phi_1]\ldots[\phi_n]$ supports ψ.

A sequence of sentences $\phi_1, \ldots, \phi_n$ entails a sentence ψ if whenever an information state is sequentially updated with $\phi_1, \ldots, \phi_n$, the resulting state is one which supports ψ, provided that along the way no free variables occur that at that point have not introduced in the referent system.

3.2 Equivalence

A suitable notion of equivalence may be expected to tell when two expressions can be substituted for each other in a meaning preserving way. Within update semantics, meaning is preserved if the update effects are. This being so, the usual definition of equivalence in terms of mutual entailment cannot be used.

For example, $\exists xPx$ and $\exists yPy$ mutually entail each other. Whenever a state s has been updated with $\exists xPx$, there will be no possibilities left

7. See van Benthem 1991a, Groenendijk and Stokhof 1991, Groenendijk and Stokhof 1990b, Veltman to appear, for some discussion of other options.

containing a world in which there are no objects that have the property P. Any such state will support $\exists y Py$. But, obviously, $\exists x Px$ and $\exists y Py$ cannot be substituted for each other in all contexts, because although they always contribute the same information about the world, they make different contributions to the discourse information.

Likewise, a characteristic feature of the dynamic entailment relation is that it allows for binding relations between quantifiers in the premises and variables in the conclusion. For example, $\exists x Px \models Px$. And it also holds that $Px \models \exists x Px$, because in any state s such that $s[Px]$ exists, it will hold that after updating s with Px, the resulting state will support $\exists x Px$. Nevertheless, the two formulae are not equivalent. Whereas an update with $\exists x \neg Px \wedge Px$ will always lead to the absurd state, replacing the second conjunct Px by $\exists x Px$ yields a consistent sequence of sentences.

On the other hand, the requirement that ϕ and ψ always have the same update effects, i.e., that $[\phi] = [\psi]$, would be too strong. Under such a definition, $\exists x \exists y Rxy$ and $\exists y \exists x Rxy$ would not be equivalent. The reason for this is that the referent system of an information state not just keeps track of which variables and pegs are present, but also of the order in which they were introduced. After updating an initial state with $\exists x \exists y Rxy$, the first peg will be associated with x, and the second with y, and only possibilities are left in which the first peg stands into the relation R to the second. In case that same state is updated with $\exists y \exists x Rxy$, things are the other way around. Still, in terms of the possible values of the variables x and y, things are exactly the same in both cases. Although $s[\exists x \exists y Rxy]$ and $s[\exists y \exists x Rxy]$ are not the same, both states will allow and support exactly the same formulae.

Similarly, updating a state with $\exists x Px$ and with $\exists x Px \wedge \exists x Px$ does not result in the same state. Consider the minimal state. After updating it with $\exists x Px$ there will be only one peg present, associated with the variable x. If it is updated with $\exists x Px \wedge \exists x Px$, there will be two pegs, and only the second is associated with the variable x, the first is no longer associated with a variable anymore. But, again, in terms of the possible values of the variable x, things are the same, and both resulting states allow and support the same formulae.

In view of this, in order for two formulae to be equivalent, it is not required that an update with either one of them always leads to exactly the *same* result, but that the results are *similar*, where the notion of similarity is defined in such a way that it ignores differences between information states which are irrelevant with regard to which formulae they allow and support.

Definition 3.5 Let $i, i' \in I, i = \langle r, g, w\rangle, i' = \langle r', g', w'\rangle$, with v and v' the domain of r and r', respectively; and let $s, s' \in S$.

i. *i is similar to i'* iff $v = v'$, $w = w'$, and $\forall x \in v{:}\, g(r(x)) = g'(r'(x))$.

ii. *s is similar to s'* iff $\forall i \in s{:}\, \exists i' \in s'{:}\, i$ is similar to i', and $\forall i' \in s'{:}\, \exists i \in s{:}\, i'$ is similar to i.

The notion of similarity robs the pegs of their identity. It does not matter what they are. The only thing that matters is what hangs on them: the variables they are associated with, and the values these variables are assigned through their mediation.

Why then have pegs to begin with? The question allows for several answers. In the present context the short answer is that they are useful. Pegs make it possible to use an existential quantifier more than once, and still formulate the semantics as an *update* semantics. Without pegs, a proper dynamic semantics for the existential quantifier involves an operation of downdate, which throws away any information about possible values of the variable involved. The introduction of such downdates complicates the formulation of the basic semantic notions. Another answer is the following. Consider a language (natural language?) without variables, or just a very limited amount of them. Then an account of anaphora demands a device like that of pegs. Moreover, if a language has other means than pronouns to establish anaphoric links (such as anaphoric definite descriptions), pegs are very useful, too.[8]

Similarity is an equivalence relation.

Definition 3.6 $\phi \equiv \psi$ iff for all information states $s{:}\, s[\phi]$ is similar to $s[\psi]$.

Under this notion of equivalence, $\exists x \exists y Rxy$ and $\exists y \exists x Rxy$ are equivalent, and so are $\exists x Px$ and $\exists x Px \wedge \exists x Px$. And $\exists x Px$ and $\exists y Py$ are not equivalent, and neither are $\exists x Px$ and Px.

8. See Groenendijk et al. to appear. See also Vermeulen 1994a, chapter 3 for a discussion of the syntactic and semantic roles of variables.

4 Illustrations

4.1 Modality

Order Matters

A characteristic feature of dynamic semantics, is that it can account for the fact that order matters in discourse. Consider:

(1) It might be raining outside. [...] It isn't raining outside.
(2) It isn't raining outside. [...] *It might be raining outside.

Given the sequential interpretation of conjunction and the interpretation of the *might*-operator as a consistency test, the unacceptability of (2) is readily explained. After an information state has been updated with the information that it is not raining, it is no longer consistent with the information that it might be raining. If, as in (1), things are presented in the opposite order, there is no problem.

So, the difference between (1) and (2) is explained by the following fact:[9]

Fact 4.1 Whereas $\Diamond p \wedge \neg p$ is consistent, $\neg p \wedge \Diamond p$ is inconsistent.

Note that the dots in example (1) are important. If they are left out, or replaced by 'and', one is more or less forced to look upon (1) as a single utterance, of a single speaker, on a single occasion. But in that case, (1) intuitively is no longer acceptable. The following fact explains this:

Fact 4.2 Although consistent, $\Diamond p \wedge \neg p$ is incoherent.

An utterance of a sentence is incoherent if no *single* information state can support it. Even though multi-speaker discourses are not explicitly introduced, the semantics explains that a discourse like (1) is only acceptable if the two sentences are uttered by different speakers in different information states, or by one and the speaker who has gained additional information in between uttering the two sentences. Only if the two sentences are taken separately, each of them can be coherent. And when the order in which they are presented is as in (1) the sequence of the two sentences is also consistent, whereas in the order presented in (2) it is not.

---

9. Propositional variables are zero-place predicates. Definition 2.5 guarantees that if p is a zero-place predicate: $w(p) \in \{0, 1\}$.

Idempotency

Another way to look at the consistency and incoherence of $\Diamond p \wedge \neg p$ is as follows. Since $\Diamond p \wedge \neg p$ is consistent, there are states that can be updated with it. But once a hearer has updated her information state with $\Diamond p \wedge \neg p$, she cannot confirm what was said. For any non-absurd state s, $s[\Diamond p \wedge \neg p]$ does not support $\Diamond p \wedge \neg p$. This means that $\Diamond p \wedge \neg p$ is not idempotent:

Fact 4.3 $\Diamond p \wedge \neg p \not\models \Diamond p \wedge \neg p$.

And this, in turn, means that dynamic entailment does not have the property of idempotency.

The reason behind this is the *non-persistence* of formulae of the form $\Diamond\phi$: a state s may support $\Diamond\phi$, whereas a more informative state s' may be inconsistent with it. If the consistency test $\Diamond p$ succeeds in a situation s, and a subsequent update with $\neg p$ succeeds too, the state $s[\Diamond p \wedge \neg p]$ is a real extension of s. The information that $s[\Diamond p \wedge \neg p]$ contains in addition to the information contained by s is the reason why $\Diamond p$ is no longer consistent with $s[\Diamond p \wedge \neg p]$.

The non-persistence of modal formulae also causes non-monotonicity of entailment:

Fact 4.4 $\Diamond p \models \Diamond p$ but $\Diamond p, \neg p \not\models \Diamond p$.

Commutativity, idempotency, and monotonicity also fail for reasons having to do with coreference rather than with modality. For example, whereas $\neg Px \wedge \exists x Px$ is consistent, $\exists x Px \wedge \neg Px$ is not. And notice that $\neg Px \wedge \exists x Px$ is not idempotent. Finally, although $\exists x Px \models Px$, it holds that $\exists x Px, \exists x \neg Px \not\models Px$.

Modality and Information

If one is told that it might be raining, this may very well constitute *real* information, on the basis of which one could decide, for example, to take an umbrella when going out. In many cases, a sentence of the form *might*-ϕ will have the effect that one becomes aware of the possibility of ϕ. The present framework is one in which possible worlds are total objects, and in which growth of information about the world is explicated in terms of elimination of possibilities. Becoming aware of a possibility can not be accounted for in a natural fashion in such an eliminative approach. It would amount to extending *partial* possibilities, rather than eliminating total ones. To account for that aspect of the meaning of *might* a constructive approach seems to be called for.

The present semantics merely takes into account that upon hearing *might*-ϕ, one checks whether one's information allows for the possibility that ϕ. This does explain the following observation. Suppose again that one is told that it might be raining. And suppose furthermore that the information one has tells that this is not so. Then, in all likelihood, one will not accept the remark just like that, one will start arguing. ('No! Look outside! The sun is shining!'). It is this aspect of the meaning of *might* that is accounted for by the semantics that was given above.

If ϕ is inconsistent with an information state s, updating s with $\Diamond\phi$ would result in $\emptyset$. But one does not want to end up in the absurd state. Hence, if that threatens to happen, one does *not* just update one's information state, but one will start arguing with whoever tries to tell that $\Diamond\phi$.

One way to look upon this, is that epistemic modal statements such as *might*-ϕ are not primarily meant as providing information about the world as such; they rather provide information about the information of the speaker. If a speaker utters *might*-ϕ the hearer may infer, on the assumption that the speaker's utterance is correct, that his information is consistent with ϕ. Since this type of higher-order information is left out of consideration here, this kind of update effect is not accounted for.

Another possible update aspect of epistemic modal statements is the following: In some situations *might*-ϕ draws attention to a hypothetical possible extension of one's information. Often this is done with the intention of saying something more about 'what if'. An example is the following sequence of sentences:

(3) It might rain. It would ruin your blue suede shoes.[10]

The effect of updating one's information state with this sequence of sentences should roughly be that it is extended with the conditional that if it rains the blue suede shoes one is wearing will be ruined, which could be a real update, and not just a consistency test.

This phenomenon, which is known under the name of *modal subordination*, is a central feature of the meaning of natural language modalities. (And, actually, it can be used as a first rate argument in favor of a dynamic treatment of it.) Nevertheless, except for a slightly more elaborated discussion at the end of the next section, it is largely ignored in the present

10. Just like stepping on them would.

paper.

For the moment it suffices to indicate that there is more to *might* than the semantics presented in this paper for the modal operator $\Diamond$ covers. On the other hand, the observations made above may have made clear that consistency testing is indeed an essential ingredient of its meaning.[11]

4.2 Coreference and Modality

Coreference

It is a characteristic feature of dynamic semantics that an existential quantifier can bind variables outside its scope. The variable in the second conjunct of (4) is bound by the quantifier in the first conjunct:

(4) $\exists x Px \wedge Qx$

A man is walking in the rain. He wears blue suede shoes.

Let n be the number of pegs in an information state s. First, s is updated with $\exists x Px$. Each possibility $\langle r, g, w\rangle \in s$ will have as many possibilities $\langle r[x/n], g[n/d], w\rangle$ as its descendants in $s[\exists x Px]$ as there are objects $d \in D$ such that $d \in w(P)$. From those, the update with Qx eliminates the possibilities i in which $i(x) \notin i(Q)$.

Exactly the same result is obtained when s is updated with:

(5) $\exists x(Px \wedge Qx)$

There is a man walking in the rain who wears blue suede shoes.

This equivalence is a basic fact of dynamic predicate logic:

Fact 4.5 $\exists x Px \wedge Qx \equiv \exists x(Px \wedge Qx)$.

With the aid of the extended binding power of the existential quantifier a compositional and incremental account of cross-sentential anaphora can be given, and the same holds for donkey-anaphora, as is guaranteed by the following equivalence:

11. Consistency testing is an essential ingredient of *might* as an *epistemic* modality. In the present paper only epistemic modalities are treated. Although the information based nature of dynamic semantics may suggest otherwise, this is not a principled limitation. Alethic modalities can be added, making it possible to implement the Kripkean distinction between metaphysical and epistemic necessity. For this purpose a set of metaphysically possible worlds must be added to each possibility. Different possibilities may contain different alternative sets of such worlds. In this way, one can account for the learnability of what is metaphysically possible, necessary, and impossible.

Fact 4.6 $\exists x Px \to Qx \equiv \forall x(Px \to Qx)$.

These equivalences are the trade mark of dynamic predicate logic. They make it possible to translate the sentences in (6) and (7) into logical formulae which reflect the structure of these sentences more transparantly than their usual logical translation (8)

(6) If a farmer owns a donkey he beats it.
$(\exists x Px \wedge \exists y(Qy \wedge Rxy)) \to Sxy$

(7) Every farmer who owns a donkey he beats it.
$\forall x(Px \wedge \exists y(Qy \wedge Rxy)) \to Sxy)$

(8) $\forall x \forall y((Px \wedge Qy \wedge Rxy) \to Sxy)$

The equivalences stated above guarantee that the formulae in (6), (7) and (8) are logically equivalent.[12]

Modality and Coreference

Modal operators are transparent to the extended binding force of existential quantifiers. In (9), the occurrence of the variable within the scope of the *might*-operator is bound by the quantifier in the first conjunct:

(9) $\exists x Px \wedge \Diamond Qx$

In this case, the second conjunct only tests whether there is at least one possibility $i \in s[\exists x Px]$, such that $i(x) \in i(Q)$. If so, the test returns the state that resulted from updating with $\exists x Px$, if not, it gives the absurd state. In particular, this means that there may be possibilities $i \in s[\exists x Px \wedge \Diamond Qx]$ such that for no $i' \in s[\exists x Px \wedge \Diamond Qx]$ it holds that $i(x) \in i'(Q)$. In other words, among the possible values for x there may be objects d that do not have the property Q in any of the worlds compatible with the information.

As is to be expected, both (10) and (11) are inconsistent:

(10) $\exists x Px \wedge \Diamond \neg Px$

(11) $\exists x Px \wedge \Diamond \forall y \neg Py$

This corresponds to the problematic nature of the following discourses:

(12) There is someone hiding in the closet. [...] He might not be hiding in the closet.

12. For an extensive discussion of the analysis of donkey-anaphora in a dynamic setting, and a comparison with their treatment in the framework of discourse representation theory, see Groenendijk and Stokhof 1991.

(13) There is someone hiding in the closet. [...] It might be that no-one is hiding in the closet.

The first discourse is unacceptable if the pronoun in the second sentence is interpreted as anaphorically linked to the indefinite in the first. The second discourse is unacceptable unless the second sentence is looked upon as a revision or a correction of the information provided by the first sentence.

Unlike (10) and (11) the following formula is *not* inconsistent:

(14) $\exists x Px \wedge \forall y \Diamond \neg Py$

The typical situation in which an information state supports (14), is where it is known that someone is hiding in the closet, but where any information on who it is, is lacking. That means that for any of the persons involved, it is consistent with the information that it is not that person who is hiding in the closet.

More formally, suppose the domain consists of just two objects, and that according to some information state just one of them has the property P, but that it does not decide which one it is. Then for each of these objects it holds that it might not have the property P.

Unlike (14), (15) *is* inconsistent:

(15) $\exists x(Px \wedge \forall y \Diamond \neg Py)$

The brackets make a difference. In updating a state s with (15), some object d is chosen, and $s[x/d][Px \wedge \forall y \Diamond \neg Py]$ is performed. In all possibilities that remain after updating $s[x/d]$ with Px, d has the property P. But then $\forall y \Diamond \neg Py$ will be inconsistent with $s[x/d][Px]$. And this holds for each choice of d. Hence (15) is inconsistent.

The fact that (14) is consistent, whereas (15) is not, means that dynamic modal predicate logic lacks some features which characterize dynamic predicate logic. It is no longer the case that $\exists x \phi \wedge \psi$ and $\exists x(\phi \wedge \psi)$ are always equivalent. This point may be elaborated.

The Case of the Broken Vase

Imagine the following situation. You and your spouse have three sons. One of them broke a vase. Your spouse is very anxious to find out who did it. Both you and your spouse know that your eldest didn't do it, he was playing outside when it must have happened. Actually, you are not interested in the question who broke the vase. But you are looking for your eldest son to help you do the dishes. He might be hiding somewhere.

In search for the culprit, your spouse has gone upstairs. Suppose your spouse hears a noise coming from the closet. If it is the shuffling of feet, your spouse will know that someone is hiding in there, but will not be able to exclude any of your three sons. In that case your spouse could utter:

(16) There is someone hiding in the closet. He might be guilty.
$\exists x Qx \wedge \Diamond Px$

But the information state of your spouse would not support:[13]

(17) There is someone hiding in the closet who might be guilty.
$\exists x(Qx \wedge \Diamond Px)$

If the situation is slightly changed, and it is imagined that the noise your spouse hears is a high-pitched voice, things are different. Now, your spouse knows it can not be your eldest, he already has a frog in his throat. In that case your spouse *can* say (17).

This also means that if your spouse yells (17) from upstairs, you can stay were you are, but if it is (16), you might run upstairs to check whether it is perhaps your aid that is hiding there.

So, there is a difference between (16) and (17),[14] and the semantics accounts for it:

Fact 4.7 $\exists x Px \wedge \Diamond Qx \not\equiv \exists x(Px \wedge \Diamond Qx)$.

For $\exists x Px \wedge \Diamond Qx$ to be supported by an information state s, it is sufficient that in any possibility in s the denotation of P is not empty, and that there is at least one possibility in which an object which has the property P also has the property Q. In particular, this leaves open the option that there is some possibility such that the object(s) satisfying P in that possibility fail to have the property Q in any possibility. This is why in the example, your spouse can correctly utter (16) in case it is possible according to the information of your spouse that your eldest is hiding there, whereas your spouse knows that he cannot be guilty of breaking the vase.

If a state s supports $\exists x(Px \wedge \Diamond Qx)$, the following holds. In any possibility in s the denotation of P is not empty. Moreover, at least one of the objects which in a possibility satisfy P must satisfy Q in some possibility. This excludes the option that there is some possibility such that the object(s) satisfying P in that possibility fail to have the property Q in any possibility. This is why your spouse can only correctly utter (17) in case the

13. Note that the interpretation of (17) that is relevant here, is marked by a specific intonation contour, which has stress on 'closet'. If the stress is on 'might', a different interpretation results, which is the same as that of (16), or so it seems.
14. Thanks to David Beaver for pointing this out.

information state of your spouse does not allow for the possibility that your eldest is hiding in the closet.

A similar observation applies to the following pair of examples.

(18) If there is someone hiding in the closet, he might be guilty.
$\exists x Qx \rightarrow \Diamond Px$

(19) Whoever is hiding in the closet might be guilty.
$\forall x(Qx \rightarrow \Diamond Px)$

Take the same situation again. Only in case your spouse heard some high-pitched voice, (19) is a correct utterance. In the other case, (19) is not supported by the information state of your spouse, and only (18) is left.

Fact 4.8 $\exists x Qx \rightarrow \Diamond Px \not\equiv \forall x(Qx \rightarrow \Diamond Px)$.

These facts are significant for at least two reasons. First, unlike in the predicate logical fragment of the language, in the full language it makes a difference whether a bound variable is inside or outside the scope of the quantifier that binds it. Secondly, since in any static semantics a variable can only be bound by a quantifier if it is inside its scope, it can never account for such differences.

Two features

There are two features of the proposed semantics which together are responsible for this result. The first is that the consistency test performed by the *might*-operator not only checks whether after an update with the formula following the *might*-operator there will be any worlds left, but also whether there will be any assignments left. Thus, even in a situation in which knowledge of the world is complete (or irrelevant), epistemic qualification of a statement may still make sense. Example:

(20) $\exists x(x^2 > 4) \wedge \Diamond(x > 2) \wedge \Diamond(x < -2)$

Consider the world that results when the the operations and relations mentioned in (20) are given their standard interpretation in the domain of real numbers. In that case (20) will be supported by any state consisting of possibilities in which only this world figures.

The second feature is that existential quantification is *not* interpreted in terms of *global* (re-)assignment. Global reassignment, which would give wrong results, reads as follows:

$$s[\exists x\phi] = (\cup_{d \in D} s[x/d])[\phi]$$

Updating with $\exists x \Diamond Px$ would output *every* $d \in D$ as a possible value for x, as long as there is *some* d that in some world compatible with the information has the property P. The present definition reads:

$$s[\exists x\phi] = \cup_{d \in D}(s[x/d][\phi])$$

Updating with $\exists x \Diamond Px$ outputs as possible values of x only those d such that in some w compatible with the information in s, d has the property P in w. If $\Diamond Px$ is within the scope of $\exists x$, the consistency test is performed one by one for each $d \in D$, and those d are eliminated as possible values for x for which the test fails.[15]

Modalities *de Dicto* and *de Re*

Note that there is a difference between $\exists x \Diamond Px$ and $\Diamond \exists x Px$. Like negation, modal operators block the binding of quantifiers inside their scope. An update of a state s with $\Diamond \exists x Px$ only tests whether there is some possibility in s in which the denotation of P is non-empty. If so, $s[\Diamond \exists x Px] = s$, if not $s[\Diamond \exists x Px] = \emptyset$. In calculating $s[\Diamond \exists x Px]$, a hypothetical update of s with $\exists x Px$ is performed. But apart from the question whether $s[\exists x Px]$ leads to the absurd state or not, the effects of this hypothetical update are ignored. In particular, the extension of the referent system of s with a new peg associated with x is not inherited by the update of s with $\Diamond \exists x Px$ as a whole. This is why the quantifier inside the scope of the modal operator has no binding force outside the scope of that operator.

In case the consistency test $\Diamond \exists x Px$ fails, not only $\Diamond \exists x Px$, but also $\exists x \Diamond Px$ leads to the absurd state. In case the test $\Diamond \exists x Px$ succeeds in a state s, $s[\exists x \Diamond Px]$ is a real extension of $s[\Diamond \exists x Px] = s$. But the additional information only concerns discourse information: s subsists in $s[\exists x \Diamond \phi]$. Although $\Diamond \exists x Px$ and $\exists x \Diamond Px$ are not equivalent, they do entail each other.[16]

The unacceptability of the discourse in (21) squares with the fact that the existential quantifier inside the scope of $\Diamond$ in the first conjunct of the formula in (21), does not bind the variable in the second conjunct:

(21) It might be the case that someone is hiding in the closet. *He broke the vase.

---

15. It is these two features which distinguish the present system from the one defined in van Eijck and Cepparello to appear.

16. Of course, when the assumption that the language users know which objects constitute the domain of discourse is dropped, $\Diamond \exists x Px$ and $\exists x \Diamond Px$ are not only non-equivalent, but it also does not hold anymore that $\Diamond \exists x Px$ entails $\exists x \Diamond \phi$.

$\Diamond\exists xPx \wedge Qx$

However, there are also cases that seem to point in a different direction:

(22) It might be the case that someone is hiding in the closet. It might be that he broke the vase.

$\Diamond\exists xPx \wedge \Diamond Qx$

Here, it seems, the pronoun in the second sentence can be interpreted as anaphorically linked with the indefinite in the first sentence. However, in the present semantics, the variable in the second conjunct is not bound by the quantifier in the first.

Modal Subordination

There is a notable difference between the discourses in (21) and (22): whereas the second sentence in (21) is in the indicative mood, the second sentence in (22), like the first, is a modal statement. The discourse in (22) is a typical example of what Roberts called 'modal subordination'.[17] The possibility of anaphoric linking between an indefinite embedded in a modal statement and a pronoun in a subsequent sentence, seems to be restricted to cases where the latter has a similar modal force.

As it stands, the analysis of modality and coreference presented here, does not account for the phenomenon of modal subordination. But it is not too difficult to see in which way the framework could be extended to be able to deal with it. Consider the well-known example:

(23) A wolf might come in. It would eat you first.

Intuitively, what this little horror story tries to tell the hearer is, first of all, that there are possibilities in which a wolf comes in; and, secondly, that in every such possibility she is the first to be consumed by the ferocious animal. In effect, the second sentence provides the information also conveyed by 'If a wolf comes in, it will eat you first'.

The first sentence is of the form $\Diamond\exists xPx$. What is needed to interpret the second sentence is to keep track of the hypothetical update of the original state s with $\exists xPx$, which is involved in the consistency test constituted by the first sentence. The second sentence is of the form *would* Qx, where its mood indicates that, in addition to the original state s, the hypothetical state $s[\exists xPx]$ must be taken into consideration. Updating with the second sentence involves a further hypothetical update of $s[\exists xPx]$ with Qx. Finally, the interpretation of the modal operator *would* effectuates the elimination

17. See Roberts 1987; Roberts 1989 for extensive discussion.

of those possibilities in s which subsist in $s[\exists xPx]$, but the descendants of which do not subsist in $s[\exists xPx][Qx]$. So, the end result can be a real update of the original state, where the less frightening possibilities where a wolf comes in and eats someone else first, or even better, eats no one at all, are eliminated. (But look at it from the bright side: The best of all possible worlds in which no wolf comes in are not eliminated!)

To make an analysis like this work the framework needs to be extended in such a way that within the update procedures intermediate hypothetical states are rememberded, rather than immediately forgotten.[18] Roughly speaking, if the next sentence is in the indicative mood, such hypothetical states can be removed from memory; if the next sentence is a modal statement, this signals that if such hypothetical states are in memory, they can be put to use, where the particular modality involved, determines the way in which they should be used.

Once the framework has been extended along these lines, not only the discourse in (23) can be handled, but also the earlier example (22) poses no problem anymore. In interpreting $\Diamond\exists xPx \wedge \Diamond Qx$ in a state s, the consistency test $\Diamond Qx$ could be performed with respect to the hypothetical state $s[\exists xPx]$. The overall effect would be the same as in case of performing the single consistency test $\Diamond(\exists xPx \wedge Qx)$. (Of course, $\Diamond\phi \wedge \Diamond\psi$ in general would still not be equivalent with $\Diamond(\phi \wedge \psi)$.)

Note that once modal subordination has been accommodated, the difference between $\Diamond\exists xPx$ and $\exists x\Diamond Px$ becomes even more apparent. Consider the following examples:

(24) It might be the case that there is someone hiding in the closet. *But he might also not be hiding in the closet.

$\Diamond\exists xPx \wedge \Diamond\neg Px$

(25) There is someone who might be hiding in the closet. But he might also not be hiding in the closet.

$\exists x\Diamond Px \wedge \Diamond\neg Px$

In case of (24), the pronoun in the second sentence can only be interpreted as anaphorically linked to the indefinite under the scope of *might* in the

18. An extension along these lines is presented in Zeevat 1992, where it is applied in an analysis of presupposition accommodation. Non-local accommodation of presuppositions also requires keeping track of (part of) the update history. Furthermore, it seems worthwhile to investigate whether, as an alternative to the analyses provided in Groenendijk and Stokhof 1990a; Dekker 1993b, anaphoric binding across negation and other 'externally static' operators, can be accounted for in this manner as well.

first sentence, if one interprets the second sentence as modally subordinated to the first. But as was pointed out above, under such an interpretation $\Diamond\exists xPx \wedge \Diamond\neg Px$ means the same as $\Diamond(\exists xPx \wedge \neg Px)$. The obvious inconsistency of the latter explains the unacceptability of the discourse in (24).

As for (25), since the indefinite has scope over *might*, the pronoun can be anaphorically linked to it without taking recourse to modal subordinationt. (In fact, because the modal operator is inside the scope of the existential quantifier, it seems that the first sentence of (25) does not count as a modal statement; and hence, the possibility of modal subordination does not even arise.) And, to be sure, the semantics as it is presented here renders (25) coherent and consistent, thus explaining the acceptability of the discourse in (25). Any state in which there is at least one possibility in which there is some object d that has the property P and some possibility in which this object d does not have the property P supports $\exists x\Diamond Px \wedge \Diamond\neg Px$.

5 Identity and Identification

Consider the following example:

(26) Someone has done it. It might be Alfred. It might not be Alfred.

$\exists! xPx \wedge \Diamond(x = a) \wedge \Diamond(x \neq a)$

$\exists x!Px$ is used as an abbreviation for $\exists xPx \wedge \forall y(Py \rightarrow y = x)$. $\exists x!Px$ expresses that there is precisely one object with the property P.

The sequence of sentences in (26) is coherent, and hence consistent. If it is continued with (27), everything remains consistent. But viewed as a single utterance, (26) followed by (27) would be incoherent.

(27) It is not Alfred. It is Bill.

$(x \neq a) \wedge (x = b)$

There are several situations in which (26) can be coherently asserted. One is the situation in which the speaker is acquainted with the person who did it, but does not know his name — his name might be Alfred, his name might not be Alfred. In this case the question who did it is decided in the information state of the speaker: in every possibility, the denotation of the predicate P is the same single object from the domain of discourse. What is not decided in the information state of the speaker is the denotation of the name a. But if the information state supports (26), then in at least one possibility the denotation of a is the object that according to the information of the speaker constitutes the denotation of P.

However, also the opposite case, in which the speaker does know perfectly well who is called Alfred, is possible. In that case the sentence reports that the question is still open whether or not this person did it. A typical example of a situation like this, not involving a name but a deictic pronoun, is this:

(28) Someone has done it. It might be you. But it might also not be you.

$\exists! xPx \wedge \Diamond(x = you) \wedge \Diamond(x \neq you)$

This is consistent and coherent, as the hearer probably would like it to be.

In ordinary modal predicate logic, the means for an adequate representation of the discourse in (26) or (28) are lacking. In ordinary (modal) predicate logic variables can not be bound unless they are in the scope of a quantifier. So, one would have to add brackets to achieve the required binding:

(29) $\exists x(Px \wedge \Diamond(x = a) \wedge \Diamond(x \neq a))$

But in case a denotes the same object in every possible world, (29) is rendered inconsistent by any modal system, including the present one.

The examples (26) and (29) show once more that whether a bound variable occurs inside or outside the syntactic scope of the quantifier that binds it can make a crucial difference. The latter type of binding is a typical feature of dynamic semantics. So, it seems that an account of the consistency of (28) calls for a dynamic set-up.[19]

5.1 Identification and Identifiers

Consider the following example:

(30) $\exists! xPx \wedge \forall y \Diamond(x = y)$

Someone has done it. It might be anyone.

Intuitively, (30) is an acceptable discourse. Formally, (30) is coherent and consistent. An information state supports (30), if it is known that just one object has the property P, but one has no idea which object it is.

Actually, under the assumption that the domain consists of more than one object, (30) entails $\forall y \Diamond(x \neq y)$. There is no static system in which $\forall y \Diamond(x \neq y)$ is consistent. In the present dynamic system it is.

Like (30), (31) expresses an ultimate form of non-identification:

(31) $\forall x \Diamond(x = a) \wedge \forall x \Diamond(x \neq a)$

19. But note that such a set-up is at best necessary, but not sufficient. For example, the system of dynamic modal predicate logic proposed in van Eijck and Cepparello to appear renders (28) inconsistent.

Anyone might be Alfred. Anyone might not be Alfred.

If an information state supports (31), it is not known of which object a is the name.

Sometimes more information is available.

Definition 5.1 Let α be a term, $s \in S$.

i. α is an *identifier in s* iff $\forall i, i' \in s\colon i(\alpha) = i'(\alpha)$.

ii. α is an *identifier* iff $\forall s\colon \alpha$ is an identifier in s.

If a term α is an identifier in an information state s, then s contains the information which object α denotes (or, who α is, in at least some sense of 'knowing who'[20]). If α is not an identifier in s, then there is at least some doubt about which object α refers to.

A term is an identifier *per se* if no matter what the information state is, it cannot fail to decide what the denotation of the term is.

Whether or not a term is an identifier in an information state can be tested:

Fact 5.1 Let α be a term which differs from x.

i. α is an identifier in s iff s supports $\forall x(\Diamond(x = \alpha) \to (x = \alpha))$.

ii. α is an identifier iff $\models \forall x(\Diamond(x = \alpha) \to (x = \alpha))$.

Identifiers are epistemically rigid designators:

Fact 5.2 Let α and β be identifiers.

i. $\models \Diamond(\alpha = \beta) \to (\alpha = \beta)$.

ii. $\models (\alpha = \beta) \to \Box(\alpha = \beta)$.

Leibniz' Law

One more aspect in which the present dynamic modal logic differs from standard ones appears if one takes a look at what happens with Leibniz' law:

Fact 5.3 If s supports $a = b$, then $s[\phi(a/x)] = s[\phi(b/x)]$

(Here, $\phi(a/x)$ is the formula which is obtained from ϕ by substituting a for free occurrences of x in ϕ.)

20. Which are many. Cf. Boër and Lycan 199; Hintikka and Hintikka 1989.

What this observation expresses is that as soon as one knows that $a = b$, $\phi(a/x)$ and $\phi(b/x)$ get the same meaning. In particular, it holds that:

$$\phi_1, \ldots, a = b, \ldots, \phi_n \models \psi(a/x) \text{ iff } \phi_1, \ldots, a = b, \ldots, \phi_n \models \psi(b/x)$$

In a standard modal semantics this holds only if a and b are rigid. In the present system it holds for all names, whether they are identifiers or not. However, since it does not generally hold that $\diamond(a = b) \models a = b$, it does not follow that:

$$\phi_1, \ldots, \diamond(a = b), \ldots, \phi_n \models \psi(a/x) \text{ iff } \phi_1, \ldots, \diamond(a = b), \ldots, \phi_n \models \psi(b/x)$$

Counterexample:

$$\diamond(a = b), \diamond(a \neq b) \not\models \diamond(b \neq b)$$

5.2 Why Identifiers are Needed

Identifiers are needed. Otherwise, if one starts from a state of ignorance, one can never really find out who is who, in the sense of coming to know the names of the objects one is talking about.

Suppose one starts out in a state of ignorance. Consider a situation with a domain of discourse consisting of just two individuals. Let a and b be names for them. Further, let there be any number of predicates. Being ignorant, one has no idea about who a is and who b is, but assume that one has learned already that $a \neq b$. Furthermore, assume that one has no idea about the denotations of the predicates.

What can one learn? A lot. For example, one can learn that Pa and $\neg Pb$; that Qa and Qb; that Rab and $\neg Rba$ and Raa and $\neg Rbb$; and so on. After having learned all this, one seems to know who has the property P: a, and no-one else. About the property Q one knows that it applies both to a and to b. Further, one has the information that a stands to both himself and to b in the relation R. Imagine that one has gained such information about all the predicates.

One's knowledge is not confined to the denotations of the predicates, one also seems to know a lot about a and b. One knows that a has the property P and the property Q, and that he stands in the relation R to himself and to b. And likewise one has learned a lot about b. Lots and lots of possibilities that one's initial state of ignorance allowed, have been eliminated. Imagine that, with respect to some fixed set of predicates and constants, one has learned anything that there is to learn in this way.

But even in that case, there are still two basic things one does not know. And because of that, there are lots of other things one does not know either. One's information still supports both $\forall x\Diamond(x = a)$, and $\forall x\Diamond(x = b)$. And that leads to a certain type of uncertainty about the predicates, too. Of the predicate P one knows that one individual has that property, but one has no idea who this is. With respect to Q things are different: since one knows that both a and b have Q, one is certain as to who is Q: everyone. As for R, there is again uncertainty. One knows which pairs form its extension, but since these are not all pairs, and since one does not know who a and b are, there is a sense in which one does not know between which individuals the relation holds.

So, although one has learned all there is to learn in this way, one has not, and will not, come to know who is a and who is b. This predicament can be formulated as follows.

Definition 5.2 Let $\langle r, g, w\rangle \in I, \langle r, g', w'\rangle \in I$.
$\langle r, g, w\rangle \simeq \langle r, g', w'\rangle$ iff there exists a bijection f from D onto D such that:

i. For every peg m in the domain of g: $g'(m) = f(g(m))$.
ii. For every individual constant a: $w'(a) = f(w(a))$.
iii. For every n-place predicate P:
$\langle d_1, \ldots, d_n\rangle \in w(P)$ iff $\langle f(d_1), \ldots, f(d_n)\rangle \in w'(P)$.

Fact 5.4 Let $\mathbf{0}$ be the minimal information state.
If $i \in \mathbf{0}[\phi_1]\ldots[\phi_n]$, then for every $i' \simeq i$, $i' \in \mathbf{0}[\phi_1]\ldots[\phi_n]$.

What this observation says is this. If one starts out from a state of ignorance — in which names are not identifiers — then, no matter how much information one gains by purely verbal means, one will never get to know to which particular object a given name refers, or which particular objects have which properties. To get this kind of information about the world, purely linguistic means are not sufficient. For identification one needs in addition non-linguistic sources of information, such as observation.

To satisfy this need, deictic demonstratives are added to the inventory of the language. It is assumed (rather naively) that if a demonstrative is used, an object is observably present in the discourse situation, which can unambiguously be pointed out to the hearer by the speaker.

Definition 5.3

i. Let $d \in D$. Then $this_d$ is a term.

ii. Let $i \in I$. Then $i(this_d) = d$.

By definition, demonstratives are identifiers. Once they are added to the language, fact 5.4 no longer holds. Expressions such as $this_d = a$ are now available, which can tell one which object a refers to.

Instantiation and Generalization

Identifiers have a special logical role, which, among others, becomes clear from their behaviour with regard to instantiation and generalization.

Suppose the domain consists of two distinct individuals d and d'. The state of total ignorance is updated with the following sentence.

(32) $a \neq b$

The resulting information state, s, supports

(33) $\Diamond(this_d = a) \wedge \Diamond(this_d = b)$

But s does not support:

(34) $\forall x \Diamond(this_d = x)$

Actually, s does not even allow (34), despite the fact that s supports the two instantiations with a and b —and these are the names of all the objects around!

The state s does support (35) and (36):

(35) $\Diamond(this_d = a) \wedge \Diamond(this_{d'} = a)$

(36) $\forall x \Diamond(x = a)$

However, at the same time the state s is inconsistent with:

(37) $\Diamond(b = a)$

which can be straightforwardly derived from (36) by universal instantiation — or so it would seem. In other words, universal instantiation is not always valid. In particular, things may go wrong if one instantiates with a term which is not known to be an identifier. Likewise, existential generalization sometimes fails:

(38) $\forall y \Diamond(y \neq a) \not\models \exists x \forall y \Diamond(y \neq x)$

Here, too, generalization is not allowed because the constant a is not an identifier.

6 Concluding Remarks

The aim of the present paper was to present an introduction to dynamic semantics, and to show that it can be fruitfully applied. Rather than by giving a bird's eye view, and advocating the approach by displaying a wide range of empirical applications, we have tried to do so by focussing on a more detailed analysis of a particular set of problems, using a particular logical language.

However, the research that has been carried out in the framework of dynamic semantics comprises both empirical studies as well as more theoretical research on a great variety of topics. In what follows, a very brief overview is given along with some references to readily accessible sources.

On the empirical side, the main focus of attention has been the analysis of pronominal co-reference, in particular donkey anaphora and intersentential anaphora of various kinds, and related problems, such as the proportion problem, modal subordination, symmetric and asymmetric quantification. Such topics are treated in Groenendijk and Stokhof 1991; Groenendijk and Stokhof 1990a; Dekker 1993b; Pagin and Westerståhl 1993; Muskens 1991. A characteristic feature of the dynamic approach is that it vindicates the traditional quantificational analysis of indefinites. This makes it possible to extend the dynamic view to a theory of dynamic generalized quantifiers (van Eijck 1993; Chierchia 1992; Blutner 1993; Kanazawa 1994). A dynamic treatment of anaphora and plurals is a closely related topic.

Other empirical phenomena that have been studied in a dynamic framework include: implicit information and scripts (Bartsch 1987); verb phrase ellipsis (Gardent 1991; van Eijck and Francez 1994); relational nouns and implicit arguments (Dekker 1993a); temporal expressions (Muskens to appear); existential sentences (Blutner 1993); epistemic modalities (Veltman to appear; Veltman et al. 1990), questions.

Other important areas of application are presuppositions (Zeevat 1992), and the analysis of default reasoning (Veltman to appear; Veltman et al. 1990).

Theoretically oriented, logical studies within the field of dynamic semantics are concerned with the formal properties of various dynamic systems. Some such studies deal with completeness, expressive power, and related topics (van Eijck and de Vries 1992; van Eijck and de Vries to appear). An algebraic view on dynamic semantics is explored in among others van Benthem 1991b; van Benthem 1991a; Visser 1994.

Other theoretical studies are directed towards: a comparison of various systems (Vermeulen 1993; Groenendijk and Stokhof 1988; Groenendijk and Stokhof 1990b; van Eijck and Cepparello to appear); a study of incrementality of contexts (Vermeulen to appear a; Vermeulen 1994b); various strategies for dealing with variables (Vermeulen to appear b); the relationship between dynamic semantics and various proof systems.

An example of a philosophical application is the analysis of the Liar paradox in a dynamic framework in Groeneveld 1994.

Acknowledgments

We owe special thanks to Paul Dekker. The present paper builds heavily on the last chapter of his thesis. His comments on various stages of the work reported here have prevented us from making many mistakes. For the remaining ones we take the blame. Maria Aloni and Jelle Gerbrandy also provided useful feed-back. Earlier versions of the paper were presented on various occasions. The first of these was the Workshop on Tense and Modality (Columbus, Ohio, July 1993). For their helpful comments, we thank the participants of that workshop, and of other events where we talked about this material. Preparation of this paper was part of the Esprit Basic Research Project Dyana (6852).

References

Bartsch, R.: 1987, Frame representations and discourse representations, *Theoretical Linguistics* 14, 65–117

Blutner, R.: 1993, Dynamic generalized quantifiers and existential sentences in natural languages, *Journal of Semantics* 10(1)

Boër, S. and Lycan, W.: 199?, *Knowing Who*, MIT Press, Cambridge: Mass.

Chierchia, G.: 1992, Anaphora and dynamic binding, *Linguistics and Philosophy* 15, 111–83

Dekker, P.: 1992, An update semantics for dynamic predicate logic, in P. Dekker and M. Stokhof (eds.), *Proceedings of the Eighth Amsterdam Colloquium*, ILLC, University of Amsterdam, Amsterdam

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

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

Gärdenfors, P.: 1984, The dynamics of belief as a basis for logic, *British Journal for The Philosophy of Sciences* 35

Gardent, C.: 1991, Dynamic semantics and vp-ellipsis, in J. van Eijck (ed.), *Logics in AI*, pp 251–267, Springer

Groenendijk, J. and Stokhof, M.: 1988, Context and information in dynamic semantics, in B. Elsendoorn and H. Bouma (eds.), *Working Models of Human Perception*, pp 457–88, Academic Press, New York

Groenendijk, J. and Stokhof, M.: 1990a, 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.: 1990b, Two theories of dynamic semantics, in J. van Eijck (ed.), *Logics in AI*, pp 55–64, Springer, Berlin

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

Groenendijk, J., Stokhof, M., and Veltman, F.: to appear, Definite descriptions in update semantics, in *Proceedings SALT V*, Austin

Groeneveld, W.: 1994, Dynamic semantics and circular propositions, *Journal of Philosophical Logic* 23(267–306)

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.

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

Hintikka, J. and Hintikka, M.: 1989, *The Logic of Epistemology and the Epistemology of Logic*, Kluwer, Dordrecht

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

Kanazawa, M.: 1994, Weak vs strong readings of donkey sentences and monotonicity inference in a dynamic setting, *Linguistics and Philosophy* 17(2), 109–148

Muskens, R.: 1991, Anaphora and the logic of change, in J. van Eijck (ed.), *Logics in AI*, pp 412–27, Springer Verlag, Berlin

Muskens, R.: to appear, Tense and the logic of change, in U. Egli et al. (eds.), *Interface Aspects of Syntax, Semantics and the Lexicon*, Benjamins, Amsterdam/Philadelphia

Pagin, P. and Westerståhl, D.: 1993, Predicate logic with flexibly binding operators and natural language semantics, *Journal of Logic, Language and Information* 2(2), 89–128

Roberts, C.: 1987, *Modal Subordination, Anaphora and Distributivity*, Ph.D. thesis, University of Massachusetts, Amherst

Roberts, C.: 1989, Modal subordination and pronominal anaphora in discourse, *Linguistics and Philosophy* 12(4), 683–722

Seuren, P. A.: 1985, *Discourse Semantics*, Blackwell, Oxford

Stalnaker, R.: 1974, Pragmatic presuppositions, in M. Munitz and P. Unger (eds.), *Semantics and Philosophy*, New York University Press, New York

van Benthem, J.: 1991a, General dynamics, *Theoretical Linguistics* 17, 159–201

van Benthem, J.: 1991b, *Language in Action*, Studies in Logic, vol. 130, North-Holland, Amsterdam

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

van Eijck, J. and Cepparello, G.: to appear, Dynamic modal predicate logic, in M. Kanazawa and C. Piñon (eds.), *Dynamics, Polarity and Quantification*, CSLI, Stanford

van Eijck, J. and de Vries, F.-J.: 1992, Dynamic interpretation and Hoare deduction, *Journal of Logic, Language and Information* 1(1), 1–44

van Eijck, J. and de Vries, F.-J.: to appear, Reasoning about update logic, *Journal of Philosophical Logic*

van Eijck, J. and Francez, N.: 1994, Verb phrase ellipsis in dynamic semantics, in M. Masuch and L. Pólós (eds.), *Applied Logic: How, What and Why?*, Kluwer, Dordrecht

Veltman, F.: to appear, Defaults in update semantics, *Journal of Philosophical Logic*

Veltman, F., Klein, E., and Moens, M.: 1990, Default reasoning and dynamic interpretation of natural language, in *Esprit '90. Proceedings of the Annual ESPRIT Conference 1990*, Kluwer, Dordrecht

Vermeulen, C.: 1994a, *Explorations of the Dynamic Environment*, Ph.D. thesis, OTS/Utrecht University, Utrecht

Vermeulen, C. F.: 1993, Sequence semantics for dynamic predicate logic, *Journal of Logic, Language and Information* 2, 217–54

Vermeulen, C. F.: 1994b, Update semantics for propositional texts, in M. Masuch and L. Pólós (eds.), *Applied Logic: How, What and Why?*, Kluwer, Dordrecht

Vermeulen, C. F.: to appear a, Incremental semantics for propositional texts, *Notre Dame Journal of Formal Logic*

Vermeulen, C. F.: to appear b, Merging without mystery. Variables in dynamic semantics, *Journal of Philosophical Logic*

Visser, A.: 1994, Actions under presuppositions, in J. van Eijck and A. Visser (eds.), *Logic and Information Flow*, MIT Press, Cambridge, Mass.

Zeevat, H.: 1992, Presupposition and accommodation in update semantics, *Journal of Semantics* 9(4), 379–412