Update Semantics for Modal Predicate Logic

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

September, 1994

1 Introduction

1.1 Static and 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 the development 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 we witness a move 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 exersized 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 we take into account the change it brings about in the hearer, more specifically in the information she has at her disposal (see Stalnaker 1979; Stalnaker 1974; Stalnaker 1984).

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; Gärdenfors 1988) 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; Kamp and Reyle 1993; Heim 1982; Heim 1983). 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; Hintikka and Kulas 1985).[1] In his original paper, Kamp describes his work explicitly as an attempt to marry the static view on meaning of 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. Cf., van Benthem and van Eijck 1982 for a comparison of game-theoretical semantics with discourse representation theory.

1.2 Dynamic semantics

Let us now turn to the sketch of one particular way of formalizing the idea of dynamic interpretation. We call it 'dynamic semantics' to distinguish it from other approaches, since, as we shall see 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 plays 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 we are concerned with is about the world, and we have 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, we can paraphrase the dynamic view as follows: 'meaning is information change potential'. Per contrast, the static view can be characterized as: 'meaning is truth conditional content'. Within a logical framework, information can be represented in terms of the parameters with respect to which interpretation of expressions is defined: assignments of values to variables, possible worlds, moments in time, and so on. Using the neutral term 'index' to refer to whatever parameters are relevant, information can be characterized in terms of indices as follows: an information state of a language user is the set of those indices which are compatible with, i.e., possible according to, the information of that user. This is a rather simple-minded approach, which needs refinement and amendment in many cases, but for the present purposes it suffices. Using this terminology we can say that the traditional static view identifies the meaning of a sentence with the set of indices in which it is true, whereas the dynamic approach takes the meaning to be a function from information states to information states. Such functions are often called 'update functions', or 'updates'.

In line with this difference, we observe 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: ϕ entails ψ iff updating an information state s with ϕ leads to an information state s' in which ψ is accepted or satisfied.[2]

The research that has been carried out in the framework of dynamic semantics comprises both empirical studies as well as more theoretical research.

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 1990; Dekker 1993b; Dekker 1994; Dekker 1988; van der Does 1994a; Pagin and Westerståhl to appear; Pagin and Westerståhl 1994; 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 1993; Kanazawa 1994; van den Berg 1994; Fernando 1994). A dynamic treatment of anaphora and plurals is a closely related topic (van den Berg 1993; van den Berg 1990; van der Does 1993).

2. See van Benthem 1989; van Benthem 1991b; van Benthem 1991a; Groenendijk and Stokhof 1990b; Veltman 1990 for some discussion of various options.

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

Other important areas of application are presuppositions (Zeevat 1992; Beaver 1993a; Beaver 1993b; Beaver 1993c; van Eijck 1994b; Krahmer 1994), and the analysis of default reasoning (Veltman 1990; 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 (Janssen 1990; van Eijck and de Vries 1992b; van Eijck and de Vries 1992a; 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. The relationship with the classical modeltheoretic approach is the subject of Benthem and Cepparello 1994.

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

A more philosophical view is developed Israel 1994, and an example of a a philosophical application is the analysis of the Liar paradox in a dynamic framework in Groeneveld 1994.

2 Pronouns and modals

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 us with 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, but we do hope to show that the dynamic perspective suggests interesting new solutions to some of them. Examples of the kind of phenomena that will be studied are provided below.

The way we will proceed is by providing 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 1990. 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 appearb.

2.1 Modals and order

That order matters in information processing is illustrated by the following simple example.

Consider the following two sequences of sentences:

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

The difference between these two sequences derives from two sources: the particular meaning of the epistemic modal expression *might*, and the relative order of the two elements. A sentence of the form *might-p* has a peculiar character. It is not so much an assertion of an (onto)logical possibility, but rather an expression of a certain epistemological condition, viz., that the possibility of p being the case is not excluded by whatever information we have at our disposal. This being so, it follows that such sentences also play a peculiar role in a dynamic perspective. Unlike an ordinary assertion p, *might-p* does not trigger update of the current information state, but rather invites the hearer to test this information state for consistency with p.[3] Now consider (1a) and (1b). In the first sequence we are first invited to test our initial information state for consistency with the information that it is raining outside. If such is the case, we stay in this information state and,

3. But see the discussion below on page 190.

hence, later on may be in a position to add the information that it does not rain. In the second sequence, however, we first update our initial information state with the information that it does not rain. And if we subsequently test this new state for consistency with the information that it does rain, we end in failure. So the order of adding information and testing for consistency does matter, a fact which can be accounted for in a dynamic semantics in a natural way.

2.2 Coreference

The following examples center around the issue of coreference of indefinites and pronominal expressions.

It is characteristic of dynamic predicate logic that its existential quantifier can bind variables outside its scope. This means that a simple sequence of sentences such as (2a) can be translated in (2b), instead of in the traditional (2c), and yet get the right interpretation:

(2) a. A man walks in the park. He wears a blue sweater.
b. $\exists x Px \wedge Qx$
c. $\exists x[Px \wedge Qx]$
d. A man wearing a blue sweater is walking in the park.

Assume some initial information state which is compatible with the information that a man is walking in the park. Processing the first conjunct of (1a) results in a state according to which there is some as yet further unspecified individual which has the property of walking in the park. The second conjunct adds the further information that this individual is wearing a blue sweater. The resulting state is one which 'contains' the information that there is a man wearing a blue sweater who is walking in the park. This information is expressed by (2d), which means that updating the initial information state with either the sequence (2a) or the sentence (2d) should result in the same state. Hence, the equivalence of (2b) and (2c), the 'natural' translations of (2a) and (2d), is imperative.

As a matter of fact, the semantic interpretation which dynamic predicate logic assigns to the existential quantifier licenses the following equivalence:

$$\exists x\phi \wedge \psi \Leftrightarrow \exists x[\phi \wedge \psi]$$

This means that the required equivalence of (2b) and (2c) is taken care of. But not only that, the use of dynamic predicate logic as a representational device has the additional advantage of providing means to build representations in a compositional, incremental manner, which is a step towards an account of the incrementality of the interpretation process itself.

If we add a *might*-operator to the language of predicate logic, we want to retain this advantage, i.e., we want to be able to translate (3a) as (3b), and get the right interpretation in this case, too:

(3) a. A man walks in the park. He might be wearing a blue sweater.
b. $\exists x Px \wedge \Diamond Qx$

Again, the existential quantifier corresponding to the indefinite term binds a variable outside its scope. But in this case the variable is inside the scope of a modal expression. And that makes a difference. Assume the same initial information state as above. Processing the first conjunct of (3a) in this state of course results again in a state according to which there is some as yet further unspecified individual which has the property of walking in the park. This time, unlike in the first case, the second conjunct does not provide new information, i.e., it does not give a further specification of this individual. Rather, what it does is invite us to check whether our information state is such that the possibility of this individual wearing a blue sweater is not excluded.

So the difference between the sequences in (2a) and (3a) is that, unlike in the former, in the latter the second conjunct does not provide new information about the individual introduced by the first. This difference also shows up in the fact that whereas (2a) is equivalent to (2d), (3a) is *not* likewise equivalent to (3d):

(3) d. A man is walking in the park who might be wearing a blue sweater

Consequently, we would not expect (3b) to be equivalent to (3c):

(3) c. $\exists x[Px \wedge \Diamond Qx]$

And that means that in the system we get when we add a *might*-operator to dynamic predicate logic we no longer can expect the equivalence referred to above to go through unconditionally.

This fact is significant also for another reason: it suggests that simple sequences such as (3a) cannot be represented in ordinary modal predicate logic for a principled reason. It is not just the case that they cannot be dealt with in an incremental fashion, as was the case with (1a) and ordinary predicate logic, the problem goes deeper: the straightforward representation (3c) simply does not give the right meaning.

A second characteristic of dynamic predicate logic is that if an existential quantifier is inside the scope of a negation, then it has no possibility to bind variables outside its scope:

(4) a. It is not the case that a man is walking in the park. *He is wearing a blue sweater.

b. $\neg\exists x Px \wedge Qx$

Just as the pronoun in the second sentence of (4a) cannot be anaphorically related to the indefinite term in the first conjunct, the variable in the second conjunct of (4b) is not bound by the existential quantifier in the first.

The *might*-operator blocks binding by existential quantifiers inside its scope of variables outside it in a similar way:

(5) a. It might be the case that a man is walking in the park. *He is wearing a blue sweater.

b. $\diamond\exists x Px \wedge Qx$

In (5a) the pronoun in the second conjunct cannot be construed as an anaphor which has the indefinite in the first conjunct as its antecedent. And, apparently, the reason is that the latter occurs inside the scope of the modal *might*.[4]

2.3 Identity and information growth

A second cluster of problems have to do with identity, especially in situations of partial information and information growth.

A simple example illustrates the kind of problem that one might run into. Let us assume that the pure demonstratives *this* and *that* are epistemically rigid designators, i.e., their reference is fixed independently of the information state. Epistemic rigidity is to be distinguished from metaphysical rigidity, which is the property of referring to the same object no matter how the world changes. A term which is metaphysically rigid need not be epistemically rigid. In fact, as long as there is the possibility of being ill-informed about its referent it will not be.[5]

That demonstratives are epistemically rigid seems a reasonable assumption, in fact it merely brings to the fore that the reference of a pure

4. Notice that if the second conjunct of (5b) would have been modal in nature as well, as, e.g., in 'He probably wears a blue sweater', anaphoric relations *are* possible. The phenomenon is known as 'modal subordination' (see Roberts 1987; Roberts 1989). Although obviously highly relevant, a formal treatment of modal subordination is beyond the scope of this paper.

5. Of course, this implies that we can not analyze epistemic possibilities simply in terms of (sets of) metaphysical possibilities, but have to devise other means to do so. This problem has been studied extensively ever since the pioneering work of Kripke, Putnam, and Donnellan on rigidity.

demonstrative is not determined by any descriptive content, but solely by the demonstration accompanying its use.[6]

Consider an information state which licenses only the following:

$$this \neq that$$

$$\neg \exists x(x \neq this \wedge x \neq that)$$

$$a \neq b$$

This is a state which carries the information that the demonstratives refer to different objects, that the names a and b also have different referents, and that there are no other objects than the two referred to demonstratively.[7] An example is a situation in which we are told that there are exactly two different solutions to a certain equation, which are pointed out to us on the blackboard: *this* one, and *that* one. Let a and b be names given to these two solutions. We can imagine that a stands for *The solution Jones found*, and b for *The solution Peters came up with.* We assume that we do not know which name refers to which solution on the blackboard. And finally, let it be given that Jones and Peters produced different results, hence we know that $a \neq b$.

It is perfectly obvious that this is a contingent information state, yet some apparently plausible arguments can be produced with would lead us to conclude to the contrary. For example, someone might reason as follows: apparently, *might* $(a = this)$ and *might* $(a = that)$; so, $\forall x$ *might* $(x = a)$; hence, *might* $(b = a)$. Or as follows: apparently, *might* $(this = a)$ and *might* $(this = b)$; so, $\forall x$ *might* $(this = x)$; hence, *might* $(this = that)$. Both arguments seem flawless, yet both lead to unacceptable conclusions.

The first argument starts with the observation that since we do not know which is solution a, the one Jones found, it might be either this one or that one. But since these are the only two possible solutions, the argument continues, it holds for any solution that it might be the one that Jones found. But then, it is concluded, also b, the solution Peters came up with, might be the one Jones found, which contradicts our assumption that they gave different answers. The problem with the first argument is the last step: apparently, universal instantiation is not licensed in this case.

---

6. This holds for pure demonstratives only, not for demonstrative expressions in general. For example, the reference of *that man* is partly determined by the demonstration accompanying *that*, partly by the descriptive content of the noun *man*.
7. Here we assume, of course, that the quantification is over some antecedently restricted domain.

The second argument proceeds differently, but to an equally unacceptable solution. From the fact that we do not know which solutions Jones and Peters produced, it is concluded that this one might be either one of them. Since their answers exhaust all possibilities, the conclusion is drawn that all solutions might be this one, from which it is derived that this solution might be identical to that one, which contradicts our initial assumptions. In this case the problem lies in the second step, the universal generalization.

What this example clearly illustrates is that when we are dealing with identities in an epistemic context, some of the normal rules of predicate logic do not go through. And we will suggest that using a dynamic perspective we will be able to explain this in an intuitive manner.

2.4 Identity and identification

Other examples of puzzling cases involving pronouns and modals are exemplified by the following formulae:

(6) $\exists x Px \wedge \Diamond a = x \wedge \Diamond a \neq x$
(7) $\exists x Px \wedge \Diamond this = x \wedge \Diamond this \neq x$
(8) $\exists x Px \wedge \Diamond \neg Px$
(9) $\exists x Px \wedge \Diamond \forall y \neg Py$
(10) $\exists x Px \wedge \forall y \Diamond Py$
(11) $\exists x Px \wedge \forall y \Diamond \neg Py$
(12) $\exists x Px \wedge \forall y \Diamond y = x \wedge \forall y \Diamond y \neq x$

Each of these sequences of formulae exemplifies a desideratum that an adequate system of modal predicate logic should meet. Starting with a suitably minimal information state, we should be able to process (sequences of sentences corresponding to) (6) and (7), and (10)–(12) and end up in a consistent information state. But (8) and (9) should not be processable in this way: the information change they induce should result in an inconsistent information state. Let us indicate very briefly what kind of puzzles are involved.

The formula in (6) would be the representation of a simple discourse such as:

(13) Someone has committed the crime. John might have done it, but then again, maybe he didn't do it.

This illustrates that modally assigning a property to an individual is 'local' in the sense that it merely states that there is some possible way in which our information state might grow that would identify John as the one who committed the crime, without thereby excluding that there is also some other way of strengthening our information which would license the claim that he is innocent. Sequences of the form of (7) show that this does not hinge on the use of proper names. Also with demonstratives, which unlike proper names are epistemically rigid, this effect appears.

For a sequence corresponding to (8) this should not be possible. Consider:

(14) Someone has committed the crime. *It might be the case that he has not done it.

Intuitively, this sequence is unacceptable: whatever our information, it cannot support both the first and the second sentence. Our first utterance introduces an individual which has the property of being the culprit. Of *this* individual we can not at the same time maintain that he might be innocent, on pain of being inconsistent. It is this process of 'introducing individuals' and predicating properties of them that an adequate analysis has to account for. In a similar manner it should be explained why (9) is out.

The pair of formulas (10) and (11) would correspond to such sequences as:

(15) Someone has committed the crime. Anyone might be him.
(16) Someone has committed the crime. Anyone might not be him.

What (10) and (11) in conjunction express is that we only have the information that someone has the property of having committed the crime, but that we have no further information as to the identity of the culprit. As far as we know it might be anyone, and at the same time the innocence of everyone individually is not in question. In a static semantics this is usually taken care of by distinguishing between quantification over (individual) concepts and quantification over concrete individuals. In the case at hand the existential quantificaton would be over concepts, whereas the universal quantifier is considered to range over individuals. A dynamic approach allows us to do away with this distinction, and handle these cases with quantification over individuals only. Formula (12) illustrates essentially the same problem, this time with identity.

In the present paper we will argue that if we take a dynamic perspective on these phenomena, they can be accounted for in an intuitive and uniform fashion. We will try to substantiate this claim in section 4. But before we can do so, we must first outline the framework that we will employ.

3 Information

3.1 Two kinds of information

In a dynamic semantics we want to explicate the meaning of a sentence as its potential to change information states. In order to be able to do so, we first of all have to specify the nature of information states. As we indicated above (see page 141) our general conception of an information state is that of a set of *possibilities*, intuitively those alternatives which are open according to our information. What the possibilities that make up information states are, depends on what we want the information to be about.

Of course, we are interested in *information about the world*. In the end, that is what counts, we want to get as good an answer to the question what the world is like as we can.[8] There are many ways in which we gather information about the world: perception, reasoning, recollection.

One particular way is through the use of language: linguistic communication. And this is what we are dealing with here: the interpretation of informative language use.[9] Such use of language is primarily focussed on answering questions about the world. But the interpretation process itself brings along its own questions. When we are engaged in a linguistic information exchange, we also have to store *discourse information*. For example, there are questions about anaphoric relations that we have to resolve. To be able to do that, we have to keep track of what we have talked about and what information we have gathered about these things; we have to maintain a model of the information of other speech participants, and so on. In the present paper we will focus on discourse information of the first kind. Discourse information of this type looks more like a book-keeping device, than real information. Yet, it is information that is essential for the interpretation of discourse, and since that is an important source of information about the world, indirectly, discourse information also provides information about the world.

The role of these two parameters of information, information about the world and discourse information, can best be illustrated by looking at an example.

---

8. Thus the charge that some (see, e.g., Bunt 1990; Kamp 1990; Israel 1994) have made against such systems as dynamic predicate logic, viz., that they are not concerned 'real' information change, we regard as unfounded (but understandable).
9. But, as we will see in section 4.5, we will always also need non-linguistic resources, to really get somewhere.

3.2 The man in the park

Suppose we have the following information about the world. We know that exactly one person is walking in the park. We do not know who it is, but we do know that it is either Alfred or Bill, and we know who they are. Furthermore, we know that Alfred wears a blue sweater, but Bill does not.

This information allows for two ways the world could possibly be like. We can represent this information by letting our information state consist of two possible worlds: a world w_1 in which Alfred walks in the park, and a world w_2 in which Bill does. In both worlds Alfred wears a blue sweater, and Bill does not.

Of course, we would like to know which of the two is the real world. We already have a partial answer to that question, in the sense that we have already discarded lots of alternatives, such as a world in which Bill is also wearing a blue sweater. However, in many cases we would strive to obtain a complete answer.

Let us start a discourse. Suppose we are told: 'A man walks in the park'. Well, this much we knew already. But still something changes. We get some discourse information. The indefinite term 'a man' introduces something that has been talked about now, and given our initial information, the further predication that this thing walks in the park turns it into something unique: it is either Alfred, in world w_1, or Bill, in world w_2.

The things that are introduced by a discourse, we call *pegs*. They are not real objects, but as we can see from the example, they are linked to real objects in each of the worlds that are still possible according to our information. Since in these worlds these real objects have all kinds of properties, and stand in all kinds of relations to each other, the association of pegs with objects in our possible worlds, tells us something about what properties the things we have been talking about have, or might have, and in which relations they stand, or might stand. In other words, it gives us information about who they might be. This information is partial as long as there is more than one possible world left in our information state, and as long as there is more than one possible value for a peg.

It does not really matter what pegs are. The only thing that counts is that we can keep them apart, and that there are enough of them, no matter how many things are introduced by the discourse. A natural choice is to use natural numbers as pegs, one by one, starting with 0.

Given that, we can represent the information state that results after processing the first sentence of our discourse in our initial information state.

We were told that a man is walking in the park. The indefinite term 'a man' introduces the first peg. And given our initial information that exactly one person is walking in the park, and that this is either Alfred or Bill, we set this peg to the individual Alfred in w_1, and to Bill in w_2. So, in the resulting information state there are two possibilities: $\{\langle 0{:}\,\mathit{Alfred}, w_1\rangle,\ \langle 0{:}\,\mathit{Bill}, w_2\rangle\}$.

Let us continue the discourse. Suppose the next thing we are told is: 'He wears a blue sweater'. Whereas the previous utterance added only to our discourse information, this utterance updates the information about the world. Assuming that 'he' refers back to the man introduced earlier, we link it to the peg 0. And what we are told is that whoever the peg 0 refers to, he wears a blue sweater. That is only possible with the value assigned to the peg 0 in w_1, which is Alfred, whom we know to be the only person around who wears a blue sweater. Hence, of the two possibilities our information allowed so far, we can eliminate one: only $\langle 0{:}\,\mathit{Alfred}, w_1\rangle$ remains.

Thus this discourse changes our information. From a state in which we did not know who is walking in the park, we are led to a state in which we do: it is Alfred, the guy who wears a blue sweater.

Note that we were not told this *literally*. Actually, the person who informed us might not even be aware of this fact herself. If, unlike us, she does *not* know who wears a blue sweater, it might be quite informative to tell her now: 'Oh, then it is Alfred'. This is an instance of the platitude 'Two heads are better than one'. The two of us, each having only partial information, can sometimes each reach complete information by just talking a bit.

Although dynamic semantics is particularly apt to model this type of information *exchange*, we will not deal with these issues here. We settle for something simpler. What we will be mainly concerned with is to model the information update of a single hearer. But every now and then we will refer to speakers, too.[10] What we will *not* model is whatever information the one speech participant may have about the information of another.[11]

Returning to our example, if the discourse continues with: 'He sees a woman', this gives rise to the introduction of a new peg, that we can set to any of the women that are around, hoping that at least one of them is seen by Alfred. Otherwise, we would loose our last possibility. Let us suppose we are lucky. Suppose we also know that Clara, a cousin of the speaker's, and Donna, her best friend, are sitting next to each other on a bench in the

10. See sections 3.10, 5 for more discussion about the hearer perspective.
11. See Dekker 1993b, chapter 5 for a first analysis in a dynamic setting.

park, and that they are seen by the man who is strolling there, which by now we know to be Alfred. Our new information state will contain again two possibilities: $\{\langle 0\colon \mathit{Alfred}; 1\colon \mathit{Clara}, w_1\rangle, \langle 0\colon \mathit{Alfred}; 1\colon \mathit{Donna}, w_1\rangle\}$.

This may seem strange. We remain equally well informed about the world, but still our information has become more partial: there are more alternatives now than there were before this step in the discourse. That is true, but notice that the added uncertainty only concerns our discourse information. And that this became more partial, is only because a new discourse question has arisen, viz., which woman the speaker might be referring to. It might be Clara, but it might also be Donna, since Alfred sees both of them. And this constitutes only a partial answer to this new discourse question, which in the end we might want to be resolved completely.

Of course, it need not be. The discourse might stop here, and we would simply forget about the issue altogether. On the other hand, the discourse might also continue, e.g., with a sentence like: 'She is my cousin'. At this point the newly introduced discourse information is needed for interpretation to be possible: we need the peg and the associated question of its identity in order to be able to process this sentence. Since we know Clara to be the cousin of the speaker, we know that it is Clara that she wants to refer to. In this particular case it does not matter that much, in this sense that we do not learn anything new about the world: we already knew that Alfred is seeing both Clara and Donna, and that Clara is the speaker's cousin. But the new discourse question is resolved, and one can easily imagine a situation where we would learn something new. And, in general, we do need discourse information to get information about the world.

The above example illustrates that information can be extended in two distinct ways: by *elimination* of possibilities, and by *extension* of the possibilities themselves. In the present set-up, information about the world extends only in the first way. This is because we view possible worlds as *total* objects, and not, say, as partial situations. With respect to discourse information both ways of extending information occur: certain values of pegs with respect to certain worlds can be eliminated during the interpretation process, and our information can be extended by the introduction of new pegs.

That discourse information can affect information about the world can be seen from the fact that if a certain value of a certain peg with respect to a certain world is eliminated, this might amount to the elimination of that world itself. This happens if this value of this peg was the only value left for this peg in this world. In our example, this happened after the second sentence had been uttered.

Elimination of possibilities in an information state amounts to getting a better, a less partial answer both to the question what the world is like, and to the discourse questions about the possible values of pegs that have already been introduced. Extending possibilities in our information state amounts to introducing new discourse questions about the possible values of newly introduced pegs. So, getting more information may consist both in getting better, i.e., more complete, answers to questions we already had, and in the addition of new questions.

3.3 Possibilities

Let us take a closer look at the possibilities that turned up in our example. Take, e.g., $\langle 0{:}\,\mathit{Alfred}; 1{:}\,\mathit{Clara}, w_1\rangle$. It consists of two things: an assignment of objects to the pegs that have been introduced, and a possible world.

Changing notation slightly, we can write the assignment as the set of pairs $\{\langle 0, \mathit{Alfred}\rangle, \langle 1, \mathit{Clara}\rangle\}$. We can look upon it as an assignment function g with the set of pegs $\{0, 1\}$ as its domain, and the objects from the domain of w_1 as its range. For simplicity, we assume that all possible worlds share the same domain D. In the end (but not in this paper), we want to give up this assumption. Assuming a single domain shared by all worlds amounts to assuming that there is complete information about what constitutes the domain (but not, of course, about what each thing is called). It is more interesting, and more realistic, to consider situations in which information about the domain is partial, too.

The assignment function forges a link between the pegs introduced by the discourse, and the world. The assignments form that part of the discourse information that connects it with information about the world. To make the link more explicit, we write a possibility as a triple $\langle n, g, w\rangle$, where n is the number of pegs introduced, g the assignment that connect the pegs to objects, and w is a possible world.

According to its set-theoretical definition, we can identify a natural number n with the set of natural numbers smaller than n. So, in the example of a possibility we were looking at, $n = 2$, i.e., $n = \{0, 1\}$. If no pegs have been introduced yet, as is the case in an initial state, we have $n = 0 = \emptyset$. In a possibility we keep track of how many different pegs have been introduced. We also keep track of the order in which they were introduced.[12] We will also call the set of pegs the *referent system* of a possibility.

12. Later on we will abstract from this feature.

The last element of a possibility is a possible world. We conceive of a possible world as a complete first order model. Since we assume that all worlds share the same domain D, a world can be identified with the interpretation function of a first order model.

Let us put these things together in a definition.

Definition 1
Let D, the *domain of discourse*, and W, the set of *possible worlds*, be two disjoint, non-empty sets.
The set of *possibilities* I based on D and W is the set of triples $i = \langle n, g, w\rangle$, where:

1. n is a natural number
2. g is a function from n into D
3. $w \in W$

We will call n the *referent system* of i

3.4 Information states

Let us now turn to information states. As was said earlier, information states are conceived of as sets of possibilities, viz., those possibilities that are still open according to our information. As we saw in the example discussed above, new pegs are introduced *globally* in an information state, i.e., they are introduced in each of its possibilities. Hence, it makes sense to restrict information states to sets of possibilities that share the same number of pegs, i.e., the same referent system. This leads to the following definition.

Definition 2
Let I be the set of possibilities based on D and W.
The set of information states S based on I is the set such that $s \in S$, iff

1. $s \subseteq I$
2. $\forall i, i' \in s$: i and i' share their referent system

The condition that all possibilities in a state share their referent system, i.e., the same (finite) number of pegs, means that $s \subset I$.

In a dynamic semantics, information states will be put to use in defining the information change potential of expressions. This means that we are primarily interested in relations between information states, more in particular when one information state can be said to *extend* (or strengthen, contain more information than) another. This is the subject of the next section.

3.5 Extending possibilities

In our introductory example, we saw already that information states can be extended in two different ways: by elimination of possibilities, and by extension of the possibilities themselves. So, we first should decide what extensions of possibilities are. We will say that one possibility extends another if each of the three elements of the two pairwise stand in an extension relation, to be defined for each of these three parameters.

Definition 3
Let $i, i' \in I, i = \langle n, g, w\rangle$ and $i' = \langle n', g', w'\rangle$.
$i \leq i'$, i' is an *extension* of i, iff

1. $n \leq n'$
2. $\forall m < n{:}\, g(m) = g'(m)$
3. $w = w'$

We will sometimes write $g \leq g'$ instead of $\forall m < n{:}\, g(m) = g'(m)$. According to this definition, for a possibility i' to be an extension of a possibility i it has to have at least the pegs i has, which moreover should be assigned the same object. Also, the world should be the same. So, one possibility extends the other if everything remains the same, except for the possible introduction of new pegs, and the assignment of values to them.

Fact 1

1. $\leq$ is a partial order on I
2. The minimal elements are $\{i \in I \mid \neg\exists i'{:}\, i' < i\} = \{\langle 0, \emptyset, w\rangle \mid w \in W\}$
3. The maximal elements are $\{i \in I \mid \neg\exists i'{:}\, i < i'\} = \emptyset$

The extension relation between possibilities is a partial order, i.e., it is reflexive, anti-symmetric and transitive. It has as many minimal elements as there are possible worlds. These are the possibilities that lack any discourse information. They can be extended by adding discourse information, i.e., by adding pegs, assigning them a value, and providing information about them. Since there infinitely many pegs, this process never comes to an end. Hence, there are no maximal elements in the extension hierarchy.

3.6 Extending information states

Having defined when one possibility extends another, we are ready to proceed with the extension relation between information states. We have seen already that information can extend in two ways: by eliminating possibilities and by extending them. This is captured by the following definition.

Definition 4
Let S be the set of information states based on I, $s, s' \in S$.
$s \leq s'$, s' is an *extension of* s, iff $\forall i' \in s' \colon \exists i \in s \colon i \leq i'$

An information state s' is an extension of an information state s if every possibility in s' is an extension of some possibility in s.

Notice that given this definition, one state being an extension of another implies nothing about the numbers of possibilities they contain. If s' is an extension of s, the former may contain *less* possibilities than the latter, but it may also consist of *more* possibilities: distinct possibilities in s' may be extensions of one possibility in s.

We encountered such a situation already at the end of the introductory example. The possibilities themselves can only extend because new discourse questions are raised about the possible values of new pegs. So, the added uncertainty can only concern new issues raised by new pegs. The questions that were already raised in s, in particular the question what the world is like, are answered at least equally well by s'.

The extension relation induces a partial ordering on the set of information states.

Fact 2

1. $\leq$ is a partial order on S
2. There is a unique state of minimal information: $\{\langle 0, \emptyset, w \rangle \mid w \in W\}$
3. There is a unique state of maximal information, for all $s \colon s \leq \emptyset$
4. $\forall s \in S \colon \forall i, i' \in s$: if $i \leq i'$, then $i = i'$.

Among the states there is a unique *state of minimal information*, which is the set of all minimal possibilities. In the state of minimal information, no pegs have been introduced yet, and all worlds are still possible. Every state is an extension of the minimal state. *Initial states*, i.e., the typical states with which a discourse starts, are (non-empty) subsets of the minimal state. In general, these will already contain information about the world, but they will not carry any discourse information. There is a unique maximal information state, $\emptyset$, which we will call *the absurd state*.

The reason that only the absurd state is maximal is that there are no maximal possibilities: one can always keep adding new questions about the values of new pegs.

That all possibilities in a state share the same pegs guarantees that all possibilities in a state are of equal rank in the extension hierarchy of possibilities. There can be no two possibilities in a state such that one is a real extension of the other.

There are states of *total information*, but these are not maximal with respect to the extension ordering. A state of total information consists of a single possibility, in particular, it contains a single world. The set of total states is $\{\{i\} \mid i \in I\}$. For any total state of information, there will always be infinitely many other non-absurd, total and non-total information states that are a real extension of it. Again, the reason is that we can always keep adding new discourse questions.[13]

The extension relation between information states being a partial order, it induces a notion of (strict) identity: $s = s'$ iff $s \leq s'$ and $s' \leq s$. This means that s and s' pose exactly the same questions, and give exactly the same (partial) answers to them. In the next two sections we will discuss reasons to also consider weaker relations of resemblance between information states.

3.7 Subsistence

The notion of one state being a real extension of another, $s < s'$, is of course: $s \leq s'$ and $s \neq s'$. One might be inclined to read this as: we are better informed in s' then we were in s. But this is not really true. For s' may be a real extension of s for two different reasons. One is that we have a better, less partial answer to the questions that were already raised in s. The other is that new discourse questions have been added to the ones we had already. If only the latter applies, then we do not have gained more information, at least not in an intuitive sense. For that to be the case, at least some possibilities that were present in s should not re-occur, be it extended or not, in s'.

To be able to express this, we need a weaker notion of identity between states (which will give rise to a stronger notion of non-identity, and hence

13. Another notion of a total information state results if we just focus on the world information: s is total with respect to the world iff $\exists w \forall i \in s{:}\, i = \langle n, g, w \rangle$.

to a stronger notion of being more informative than that of being a real extension). We call call it *subsistence*. We say that a state s subsists in s' if the questions that are raised in s, receive the *same* answer in s' as in s, at the same time allowing that in s' new discourse questions are raised.

Definition 5

Let $s, s' \in S$, $i \in s$.

1. *i subsists in s'* iff $\exists i' \in s' : i \leq i'$
2. *s subsists in s'*, $s \stackrel{\leftharpoonup}{=} s'$ iff $s \leq s'$ and $\forall i \in s : i$ subsists in s'

Not only is it required that every possibility in s' be an extension of some possibility in s, but also that every possibility in s have some extension in s'. This means that all information about the world, and about the possible values of pegs already present in s, is the same in s and in s'. The only difference that is allowed is the introduction of new pegs, and information about their values.[14]

Subsistence is a dynamic, and hence not a symmetric notion of resemblance. Not everything remains the same, but only what was. New things may come up. Now we also have the means to say that *not* everything that was, has remained: $s \leq s'$ and $s \not\stackrel{\leftharpoonup}{=} s'$.

We will see in section 3.10 that the notion of subsistence plays an essential role in defining a dynamic notion of entailment. Another notion, that of similarity of states, to discussed in the next section, will turn out important in defining a notion of equivalence.

3.8 Similarity

Another reason why strict identity of states turns out to be less useful is the following.

In an information state we do not just keep track of the number of pegs that have been introduced, but also of the order in which they were.

---

14. Note that we could have used the second condition to define a separate notion $s \sqsubseteq s'$, of s being a substate of s', providing yet another partial order on information states. If $s \sqsubseteq s'$ this means that with respect to the questions present in s, s gives at least as good an answer as s' does, and there may be new discourse questions present in s'. In a sense, we tend to better off in s than in s' if $s \sqsubseteq s'$, there tend to be less questions, and we tend to have a better answer to the questions we have than in s'.
From a dynamic point of view, it is a very unintuitive notion. We cannot read it from left to right and look upon s' as something we might arrive at after having been in s, unless, that is, we conceive it possible that our information state has been downdated, rather than updated. Still we mention the notion here, since $\stackrel{\leftharpoonup}{=}$ is the composition of $\leq$ and $\sqsubseteq$.

This means that two states s and s' which are exactly the same, except for the fact that whatever we know about peg n in s, we know about peg m in s', and vice versa, are not related to each other in terms of extension, or for that matter, in terms of subsistence.

For some purposes, keeping track of order is useful discourse information. Think of anaphoric expressions such as 'the former' and 'the latter', the antecedents of which depend on the order in which pegs have been introduced. Yet, in many other cases we want to abstract away from this aspect of discourse information.

In order to be able to do this, we introduce a notion of isomorphism between states. Two states s and s' are said to be isomorphic, if we can map pegs n in s onto pegs n' in s' in such a way that whatever information we have in s about n, we have in s' about n'.

Definition 6
Let $s, s' \in S$, and $\langle n, g, w\rangle \in s$, and $\langle n', g', w'\rangle \in s'$.
$s \simeq s'$, s and s' are *isomorphic*, iff

1. s and s' share their referent system n
2. There is a bijection f from n into n', such that:
   $\langle n, g, w\rangle \in s$ iff $\langle n, g', w\rangle \in s'$, where $\forall m \in n$: $g(m) = g'(f(m))$

Two states are isomorphic if they only differ in the order in which the values of their pegs were introduced. The possibility to abstract away from that depends on the fact that pegs are arbitrary. It does not really matter what kind of things they are, it only matters how many of these things there are in an information state.

We may even go a bit further. Suppose there are two pegs in a state which in every possibility are assigned the same value. Then we might say that these two pegs are not really different after all, since what determines their identity, i.e., what real object they could stand for in each of our possibilities, is the same. Although not formally identical, they are *indiscernible*. This is defined as follows.

Definition 7
Let $s \in S$, where n be the referent system of s, $m, m' < n$.
$m \equiv_s m'$, m and m' are *indiscernible in s* iff $\forall i = \langle n, g, w\rangle \in s$: $g(m) = g(m')$

The relation of indiscernibility is an equivalence relation. In view of that it makes sense to reduce states that contain such indiscernible pegs by packing them into one. The following definition says when in this sense one state is a (partial) reduction of another.

Definition 8
Let $s, s' \in S$, with referent systems n and n' respectively.
s is a *reduct* of s' iff
1. $s \sqsubseteq s'$
2. $\forall m$: if $m \notin n$ and $m \in n'$, then $\exists m' \in n{:}\, m \equiv_{s'} m'$

The reduction relation induces yet another partial order on the set of states.

Now we can define similarity between states by saying that two states are similar if each of them can be reduced in such a way that their reductions are isomorphic.

Definition 9
Let $s, s' \in S$.
$s \approx s'$, s and s' are *similar*, iff $\exists t{:}\, t \simeq s, \exists t'{:}\, t' \simeq s'$ such that $\exists r{:}\, r$ is a reduct of t and r is a reduct of t'

Similarity is an equivalence relation between states. It is much weaker than full identity, leaving out of consideration the order in which pegs are introduced, and disregarding indiscernible ones.

Enough about information states, let us now start doing something with them.

3.9 Updates

Information is not something static, it changes all the time. It may change in many different ways, one important source being linguistic communication. Linguistic expressions function in this process primarily through their meaning, which is intrinsically related to the change in information that a communicative act involving them brings about. This is what is captured in the slogan of dynamic semantics: 'The meaning of a sentence is its information change potential'. In terms of the framework developed here, this slogan is implemented by taking the meaning of a sentence to be a function from information states to information states.[15]

15. It could also be viewed as a relation between such states, to capture the possibility of indeterminacy. For example, (syntactic and semantic) ambiguity could conceivably be handled in this way. We will not consider this possibility here, and restrict ourselves to

Definition 10
Let S be the set of information states based on a set of possibilities I.
A *state transformer on* S is a partial function from S to S.

We use postfix notation, and write $s[\tau]$ to denote the result of transforming s by τ. The postfix notation is especially perspicuous when we write down a sequence of transformations: $s[\tau][\tau']$ denotes the result of first transforming s by τ, and next transforming the state that results from that, i.e., $s[\tau]$, by τ'.

A state transformer τ may be a *partial* function, because whether a state s can in fact be transformed to some state s' by τ may depend on the fulfillment of certain constraints. If a state s does not meet them, then τ can not be applied to s, and hence $s[\tau]$ does not exist.

This leads us to consider the following property of state transformers:

Definition 11
Let τ be a state transformer on S.
τ is *safe* iff $\forall s \in S \colon s[\tau]$ exists

If a transformer is safe, then it is a total function on the set of information states. Since the composition of two total functions is itself a total function, we have that:

Fact 3
If τ and τ' are safe, then the composition of the τ and τ' is safe

This guarantees that a sequence of transformers is safe if all the transformers in the sequence are. Notice that for a sequence to be safe it is not necessary that all the transformations in the sequence are. Sequences are ordered. If some state transformer in a sequence has certain preconditions, i.e., is not safe, one that comes before it may make sure these conditions are met, thus allowing the sequence as a whole to be safe.

A class of state transformers that has our special interest are the updates:

Definition 12

---

functions. To account for indeterminacy, it seems more convenient to take a set of outcomes as the deterministic output.

Let τ be a state transformer on S.
τ is an *update* iff $\forall s \in S$ such that $s[\tau]$ exists: $s \leq s[\tau]$

Updates always transform a state into an extension of it. Since the relation of extension between states is transitive, we have that:

Fact 4
If τ is an update and τ' is an update, then their composition is an update

This guarantees that a sequence of updates is itself an update.

One might wonder how an ordinary sentence ϕ can be an update, because it is always possible that our information state contains information that is inconsistent with ϕ, and in that case no update would result. If we trust information that we already have more than any new information, we will simply refuse to update in such a case.

But this is not a good reason for thinking that no sentence is an update. In the situation where our information is in conflict with ϕ, updating our information state with ϕ would lead to the absurd state, which, as we have seen, is an extension of every state. Note that according to the definition, this does not prevent ϕ from being an update. In fact, it helps to *explain* why you refuse to simply update with such a sentence if things are like this: you do not want to end up in the absurd state. So what you do in such a situation is that you start arguing about things, or that you first revise your information to enable it to be updated with the sentence in question without going absurd. And notice that in order even to get into the position that you can decide how you should react, you *first* have to try and interpret the sentence, i.e., update 'hypothetically' with it. For only then you know it would lead to the absurd state, in which case you may react as seems appropriate.

Hence, as far as the *interpretation* of ordinary declarative sentences is concerned, it is certainly not so strange to view them as updates. On the contrary, we think it is a global constraint on a dynamic semantics for declarative sentences that their interpretation is such that they are guaranteed to be updates.

Of course, this is not to deny that there is such a thing as deletion and revision of information, besides update. The point is merely that unlike update, deletion and revision are not *directly* involved in interpretation, at least not in that of ordinary declarative sentences.[16] Update is a core semantic notion, deletion and revision are not.

Updates extend information states, either by elimination of possibilities, or by extension of possibilities, or both. A typical class of updates extend only in the first way:

Definition 13
Let τ be a state transformer on S.
τ is an *eliminative* or *non-extending* update iff $\forall s \in S$ such that $s[\tau]$ exists: $s[\tau] \subseteq s$

Eliminative updates are eliminative in the sense that they *at most* change an input state by eliminating possibilities in it. They are non-extending in the sense that they do not result in a *real* extension of the referent system of the input state, they do not add new discourse questions. Per contrast, non-eliminative updates are extending in the sense that they always give rise to real extensions of the possibilities in the input state. A non-eliminative update *may* eliminate possibilities in this sense that for some possibility in the input state no extension of it occurs in the output state.

A simple consequence is that a sequence of eliminative updates is itself eliminative:

Fact 5
If τ and τ' are both eliminative updates, then their composition is an eliminative update

Another special category of updates that needs to be distinguished is that of tests:

Definition 14
Let τ be a state transformer on S.
τ is a *test* iff $\forall s \in S$ such that $s[\phi]$ exists: $s[\tau] \leftharpoonup s$ or $s[\tau] = \emptyset$

A state transformer which has this property tests whether an information state fulfills certain requirements. If it does, the test returns a state in which the input state subsist. If a state fails, the absurd state results. If a state

16. Of course there are constructions in natural language, such as the counterfactual, which do involve revision.

is tested successfully every possibility subsists in the output state. In that sense the test does not supply any new information, at most it introduces some new discourse questions.

For tests we observe the following:

Fact 6

1. If τ and τ' are both tests, then their composition is a test.
2. If τ is eliminative and τ is a test, then $\forall s \in S$ such that τ exists in s: $s[\phi] = s$ or $s[\phi] = \emptyset$

A sequence of tests cannot fail to be a test itself. Eliminative tests either leave the state as it is, or output the absurd state.

Another important class of state transformers are those that are distributive. Unlike the properties we introduced so far, which concerned the relation between input and output states, distributivity is a property that concerns the way in which the transformation operates. A distributive transformer operates in a pointwise manner on the possibilities in a state:

Definition 15

Let τ be a state transformer on S.
τ is a *distributive* state transformer iff $\forall s \in S$: $s[\tau] = \cup_{i \in s}\{i\}[\tau]$

An state transformer is distributive if it can be executed by applying it on the states that consist of each possibility in the input state separately, and then collecting the results. This means that a distributive state transformer is not sensitive to global properties of an information state, i.e., those properties of a state which are not also properties of its elements.

Fact 7

If τ and τ' are distributive, then there composition is distributive

A sequence of distributive transformers is itself a distributive state transformer. This means that $i \in s[\tau][\tau']$ iff $\exists j \in s$: $i \in \{j\}[\tau][\tau']$.

It may happen that a state transformer is both distributive, and an eliminative update. Such state transformers we call classical updates:

Definition 16

τ is a *classical* update iff τ is distributive and τ is an eliminative update

The reason to call these updates classical is that they are not really in need of dynamic interpretation, as we shall see shortly.

First we notice that if τ is an eliminative update, then τ is an eliminative test with respect to each total information state $\{i\}$: for all $i \in I$: $\{i\}[\tau] = \{i\}$ or $= \emptyset$. If τ is also distributive, then $s[\tau]$ is the logical sum of the test $\{i\}[\tau]$ on each $i \in s$. Thus we can identify $s[\tau]$ with $\{i \in s \mid \{i\}[\tau] \neq \emptyset\}$. More specifically, for classical updates it makes sense to define the notion of a proposition:

Definition 17
If τ is a classical update, then P_τ, the *proposition expressed by* τ, $= \{i \in I \mid \{i\}[\tau]$ exists and $\{i\}[\tau] \neq \emptyset\}$

The classical nature of classical updates can then be stated as follows:

Fact 8
For all classical updates τ: if $s[\tau]$ exists, then $s[\tau] = s \cap P_\tau$

This fact implies that if we consider only classical updates, we can capture the update effects they induce simply by taking an information state in conjunction with the proposition expressed by the classical update. But this means that if we take only classical updates in consideration, then, rather than interpreting them as state transformers, we can start out by giving a classical truth definition, and define the update effects in the global manner indicated by fact 8.

The notion of τ being true with respect to a possibility $i \in I$ is defined by the condition $\{i\}[\tau] \neq \emptyset$. Then the proposition P_τ is the set of possibilities with respect to which τ is true. The test $\{i\}[\tau]$ tests for the truth of τ in i. And the update of s with τ is $s \cap P_\tau$.

So, for classical updates there is no need to state their meaning in dynamic terms, i.e., as information change potential. A classical, static interpretation in terms of information content suffices. Conversely, it only makes sense to define meaning in terms of information change potential if at least some of the expressions we want to interpret are non-distributive or non-eliminative updates.

3.10 Support, consistency, inconsistency

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 primarily related to. Hence, the notions of truth and falsity can not be expected to occupy the same central position as they do in static semantics. More suited to the information oriented approach are such notions as support, consistency, and inconsistency.

Let us so start with the latter two. If an update of our information state with a sentence would lead to the absurd state of inconsistent information, we would reject that sentence, we would refuse to update with it. We can only allow a sentence if updating our information state with it does not lead to inconsistency. If a sentence is consistent with an information state, we say that the stae *allows* the sentence, and if it is inconsistent with it, we say that the state *forbids* the sentence:

Definition 18
Let s be an information state, ϕ a sentence.

1. s *allows* ϕ iff $s[\phi]$ exists and $s[\phi] \neq \emptyset$
2. s *forbids* ϕ iff $s[\phi] = \emptyset$

Straightforward generalizations of these notions are that a sentence is consistent *per se* if there is some state that allows it, and a sentence is inconsistent *per se* if there is no state that allows it.

Definition 19
Let ϕ be a sentence, S the set of information states.

1. ϕ is *consistent* iff $\exists s \in S$: s allows ϕ
2. ϕ is *inconsistent* iff $\forall s \in S$: if $s[\phi]$ exists, then s forbids ϕ

Judgments about the consistency and inconsistency of sentences in information states are an important source of facts for a semantic theory.

If an information state allows a sentence, and we update with it, then we will usually arrive in a new state in which we have more information than we had before. But this need not always be the case, it may be that the sentence does not tell us something we did not know already. In such a situation, we can say that our information already supports that sentence. One might think that s supporting ϕ can be defined in terms of identity: $s = s[\phi]$. But recall that identity of states is very strict. The requirement that $s[\phi]$ be the same state as s, amounts to requiring that all questions, both the question what the world is like, and the discourse questions about the possible referents of pegs, are the same and receive the same (partial) answers. But this is asking too much. It is sufficient that all questions we had already in s receive the same answers in $s[\phi]$, allowing for the introduction of new discourse questions about pegs introduced in $s[\phi]$. This requirement can be stated in terms of the notion of subsistence:

Definition 20
Let s be an information state, ϕ a sentence.
s supports ϕ iff $s[\phi]$ exists and $s \preceq s[\phi]$

In terms of support various pragmatic aspects of linguistic exchanges can be captured. For example, if a speaker is to utter a sentence correctly, then her information state should support the sentence. This is Grice's Maxim of Quality. The Maxim of Quantity requires minimally that the speaker also believe that the information state of the hearer does not already support the sentence, and that she not believe that the information state of the hearer forbids the sentence.

Generally, if a sentence is to make any sense at all, there should be at least one non-absurd state that supports it. This we call *coherence*:

Definition 21
Let ϕ be a sentence, S the set of information states.
ϕ *is coherent* iff $\exists s \in S\colon s \neq \emptyset$ and s supports ϕ

Intuitions about coherence are another source of semantic facts. From the perspective of the hearer, an update with an incoherent sentence is to be rejected, since no speaker can possibly utter such a sentence correctly.

We note the following:

Fact 9
For all ϕ such that $s[\phi]$ supports ϕ it holds that ϕ is consistent iff ϕ is coherent

For sentences that give rise to idempotent updates there is no difference between consistency and coherence. Such sentences are called 'acceptable':

Definition 22
Let ϕ be a sentence.
ϕ is *acceptable* iff ϕ is consistent and coherent

For idempotent sequences consistency or coherence is sufficient for acceptability, but not all sequences have this property. For example, as we shall argue in more detail later, a sequence of the form $\Diamond p \wedge \neg p$ is consistent, but not coherent, and hence unacceptable.

It is important to bear in mind that what we are primarily modeling is the hearer's point of view. As far as her information is concerned there is nothing wrong with a sequence such as $\Diamond p \wedge \neg p$: it is consistent, i.e., it is perfectly possible to update, first with $\Diamond p$, and next with $\neg p$. Still, the sequence is not acceptable (as one single utterance) and hence, as such, cannot be 'taken in' by the hearer, because it is not coherent: there can be no *one* information state that supports it. For if such a state supports $\neg p$, the test $\Diamond p$ would fail, and vice versa. Were we to consider such a sequence as a correct utterance nevertheless, we would need to assume that the speaker's information had changed between his utterance of the two sentences of which it consists. This might very well happen in a situation in which information from outside becomes available, e.g., through observation. Or we could interpret the sequence as a 'life' report of an ongoing series of events. Or we could accept the sequence if its constituents came from different sources, as is the case in a multi-speaker setting.

These remarks should warn the reader that when a sequence is marked unacceptable, such judgement is based on the assumption that it is a one-speaker discourse with no information change taking place on part of the speaker during his utterance.

In conclusion, we see that an information state s and a sentence ϕ can be related in various ways. First of all, the update of s with ϕ may exist, or it may not. Only in the former case do the other options apply. The first of those is that s allows ϕ, but does not yet support it. This is the ideal situation for a hearer to update with it. Secondly, s may support ϕ, which is the ideal state for a speaker to utter ϕ, if she believes that the hearer is in the previous situation. Third, s may both support and forbid ϕ. This is a rare case, which holds only if s is the absurd state. The absurd state supports everything and everything is forbidden by it, which is why it is absurd. Finally, s may neither support nor allow ϕ, in which situation a speaker cannot utter it, and a hearer cannot update with it.

3.11 Entailment and equivalence

The usual static notion of entailment is defined in terms of truth. But in view its peripheral role in dynamic semantics, it will not be of much use when we want to define a dynamic notion of entailment. Given that support is the dynamic counterpart of truth, it is to be expected that entailment is defined in terms of that notion. The basic intuition is this. If updating an information state with the sequence of sentences $\phi_1, \ldots, \phi_n$ always results in a state which supports sentence ψ, we may say that $\phi_1, \ldots, \phi_n$ entail ψ. For in such a situation ψ does not add any information that is really new. At most it adds some new discourse questions.

Definition 23
Let $\phi_1, \ldots, \phi_n, \psi$ be sentences, S the set of information states.
$\phi_1, \ldots, \phi_n \models \psi$ iff $\forall s \in S$: if $s[\phi_1] \ldots [\phi_n][\psi]$ exists, then $s[\phi_1] \ldots [\phi_n]$ supports ψ

As usual, the limiting case of there being no premises, gives us a notion of universal validity: $\models \phi$ iff ϕ is supported by every state (in which an update with it exists). Intuitions about entailment and validity, too, are an important source of semantic data.

In static semantics, the notion of equivalence is usually defined as mutual entailment. But from a dynamic perspective, this is not a very natural choice. For the dynamic notion of entailment is ordered: $\phi \models \psi$ iff whenever you *first* update with ϕ, ψ is supported *afterwards*. If we could reverse the order, it would not matter. But what we want is a notion of equivalence that respects order sensitivity, i.e., one that tells us which sentences ϕ we can always, in every step in a discourse, replace safely by which sentences ψ. This implies that ϕ and ψ should have the same update effects.

Again, as was the case with the notion of support, strict identity of update effects is usually not necessary. In particular, it seems that, except in very special cases, we will not care too much about the order in which pegs have been introduced. If two updates introduce equally many discernible pegs, on which the same things are hanging with respect to each possible world, and leave us the same possible worlds as alternatives, then it seems that they are enough alike to be considered equivalent.

Definition 24
Let ϕ and ψ be sentences, S the set of information states.
$\phi \equiv \psi$ iff $\forall s \in S$: $s[\phi] \approx s[\psi]$

Again, equivalence, just like entailment, coherence, and acceptability, is an important source of semantic facts. If we have given a semantic interpretation for a particular language, we can check it, by seeing whether its predictions concerning equivalence, entailment, coherence, and acceptability, are warranted by our semantic intuitions.

Before we can turn to the semantic interpretation of the language of modal predicate logic, we have to add one more features to our information states.

3.12 Variables

The language that we want to give a dynamic semantics for is a logical language that has variables and quantifiers. This means that we will be dealing, not with natural language expressions such as 'a man', as was the case in our example, but with their formal counterparts, i.e., expressions such as $\exists x Px$. The latter, like the former, give rise to the introduction of new pegs, but unlike the former, they invite us to keep track in the discourse information of which variable is associated with which peg.[17]

Let us start with giving an informal sketch of how we propose to deal with quantifiers and variables, since that may help to understand the nature of possibilities and information states. These will be attuned to the way quantifiers and variables affect information.

So far our possibilities consisted of a number of pegs, an assignment of values to them, and a possible world. Now we add two more things. We also keep track of the variables that are in active use, and we associate each of them with one of the pegs.

---

17. In a certain sense natural languages, which do not have explicit variables, are worse off than logical languages that has them. Having none, is more or less the same as having just one, one that gets re-introduced and re-set all the time. In a logical language with explicit variables, we have at least the possibility right at hand to use different variables, and prevent having to re-introduce and re-set.
Notice that from a certain perspective not having variables is easier: there would be no need to complicate the notion of a possibility. On the other hand it would lead to complications elsewhere. These concern mainly the way we figure out anaphoric relations. Suppose we have been reading about a man and a boy and ..., and we come across the pronoun 'he'. In such a situation there may be several ways to link the pronoun to previously introduced male objects, an obvious complication. In the everyday practice of using our language this usually does not bother us too much, since we have all kinds of tricks to help us out. But in certain situations, for example if we want to define something precisely, things can become pretty hard, and then we find ourselves using variables in natural language on an *ad hoc* basis.
The main advantage of using variables is that we can be very explicit about the intended anaphoric relations. (Cf. the usual linguistic practice of indexing and co-indexing noun phrases and pronouns to disambiguate with respect to anaphoric relations.) That having variables complicates the nature of the possibilities should not surprise us, then. It makes other things easier. The complications are *inherently* there, so they will turn up somewhere. It is just that they turn up at different spots in natural and logical languages.

If we meet a quantifier $\exists x$ we do the following with each of the possibilities in our information state:

1. We add x to the set of variables
2. We add the next peg to the set of pegs
3. We associate x with the new peg
4. We assign a suitable object to the new peg
5. What is suitable depends on the world

Two remarks are in order. First, we will allow for the possibility that x is already in the set of variables, in which case nothing is added. This is because we allow for the usual practice of re-using a quantifier. If this happens, we act as if the variable were new, and associate it with the newly introduced peg. Since the association will be done by a function, it means that the link between x and the peg it was associated with before, is destroyed.

Second, we apply this procedure uniformly on all the possibilities in a state. Pegs and variables are only introduced by quantifiers in this way. This means that starting from a state without pegs and variables, characteristically an initial state, we can be sure that each possibility in a state that a discourse can lead us to, will contain the same set of variables, the same set of pegs, and the same association of variables with pegs.

Variables, pegs, and the link between them we call a *referent system*. Each information state has a unique referent system. In effect, this amounts to the assumption that there is no uncertainty about this type of discourse information.[18] The reason to copy it in all the possibilities in a state, rather than setting it apart, is that it makes some definitions run more smoothly. What will be different in the alternative possibilities in a state is that they may assign different possible values to the pegs, and that they may contain a different possible world.

3.13 New possibilities

We now re-define the set of possibilities, taking into account that our logical language has variables:

Definition 25
Let D, the *domain of discourse*, and W, the set of *possible worlds*, be two disjoint non-empty sets.

---

18. But note that if we allow it also to contain information about other aspects of a discourse, such as syntactic and semantic structure, this assumption need no longer be warranted.

The set of *possibilities* based on D, W is the set I of quintuples $i = \langle v, n, r, g, w \rangle$, where:

1. v is a finite subset of variables
2. n is a natural number
3. r is an injection from v into n
4. g is a function from n into D
5. $w \in W$

We call $\langle v, n, r \rangle$, or r for short, the *referent system* of i.

What is new is the set v of variables that are in use, and the function r that associates the elements of v with elements of n, i.e., with pegs. That r is an injection means that there may be pegs around with which no variable is associated anymore. This will always be the result of re-using one and the same quantifier.

Since variables are associated with pegs, they are also assigned an object, indirectly, *via* the peg they are associated with. The composition of g and r assigns values to variables: $g(r(x)) \in D$.

There are two reasons for assigning values to variables in this roundabout way. The first is that variables are an artifact of logical languages. Natural language by and large does without them. The pegs are motivated independently of the variables. We need them to hang on discourse information for the interpretation of whatever language we want to deal with. Secondly, if we would leave pegs out nevertheless, and would only have variables and assignments of objects directly to them, difficulties arise when defining extension of possibilities while allowing for re-using a quantifier.[19]

Information states remain defined as sets of possibilities which share a referent system. This means that all possibilities in an information state have the same set of variables in use, the same number of pegs, and that the variables in use are associated with the same pegs. The motivation for this is as before: variables and pegs are introduced with respect to an information state.

The definition of when one possibility extends another needs to be restated. We have to add what it means for a set of variables and for an association function to extend another.

Definition 26

19. This is why in other proposals that do not introduce things like pegs, one way or another, re-using of quantifiers is ruled out.

Let $i, i' \in I, i = \langle v, n, r, g, w \rangle$ and $i' = \langle v', n', r', g', w' \rangle$.
$i \leq i'$, i' is an *extension* of i, iff

1. $v \subseteq v'$
2. (a) If $x \in v$ then $r(x) = r'(x)$ or $n \leq r'(x)$
   (b) If $x \notin v$ and $x \in v'$ then $n \leq r'(x)$
3. $\langle n, g, w \rangle \leq \langle n', g', w' \rangle$

According to the first clause no variables present in i are lost i'. The second clause says that the variables present in i, in i' either remain associated with the same peg, or are associated with a peg that was not in i. This relates to re-using a quantifier: re-used variables are treated in the same way as new variables, they are associated with a new peg. The third clause requires that new variables be associated with new pegs. The fourth clause, finally, requires that with respect to the pegs, their values and the world, i' be an extension of i in the sense defined earlier.

The notion of when one state is an extension of another remains the same, and this holds also for the notion of subsistence.

What we do have to redefine is the notion of similarity.

Definition 27

1. Let $i, i' \in I, i = \langle v, n, r, g, w \rangle, i' = \langle v', n', r', g', w' \rangle$.
   i is similar to i', $i \approx i'$ iff $v = v' \,\&\, w = w' \,\&\, \forall x \in v\colon g(r(x)) = g'(r'(x))$
2. Let $s, s' \in S$.
   s is similar to s', $s \approx s'$ iff
   (a) $\forall i \in s\colon \exists i' \in s'\colon i \approx i'$
   (b) $\forall i' \in s'\colon \exists i \in s\colon i' \approx i$

For two states to be similar, they must have the same set of variables. The variables may be associated with different pegs, but the values assigned to them via the pegs should be the same. So, the order in which variables are introduced, which is the order in which they introduce pegs, may differ. Like in the previous notion of similarity, this one, too, abstracts from order. The 'identity' of pegs does not matter, what is important is the values variables are assigned through them. Note also that pegs that are not associated with a variable, play no role whatsoever. Notice, finally, that it may happen that there are two pegs around in the states we compare which have the same value in each possibility, but which remain distinguishable, as long as they are associated with different variables.

4 Co-reference, modality, and identity

4.1 Co-reference

4.1.1 Terms and predicates

The non-logical vocabulary of our language consists of terms, i.e., individual constants $a, b, c, \ldots$ and variables $x, y, z, \ldots$; and of n-place predicates $P, Q, R \ldots$, $0 \leq n$. For 0-place predicates, i.e., propositional variables, we use $p, q, r \ldots$.[20]

The possibilities contain all that is needed to interpret the basic vocabulary. The composition of the assignment function g and the association function r assigns values to variables. The world w is conceived of as the interpretation function of a first order model, and hence it takes care of the rest.

Definition 28
Let α be a basic expression, $i = \langle v, n, r, g, w \rangle \in I$, I based upon W and D.

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

The non-existence of a variable x in a state s, i.e., when x is not in the set of variables v of the referent system of s, will be the only source of partiality of our updates.[21]

4.1.2 Atomic updates

We our now ready to define the first kind of updates, the ones associated with atomic formulae.

Definition 29

1. • If R is an n-place predicate, and $t_1 \ldots t_n$ are terms, then $(Rt_1, \ldots t_n)$ is a formula
   • $s[Rt_1 \ldots t_n] = \{i \in s \mid \langle i(t_1), \ldots, i(t_n) \rangle \in i(R)\}$
2. • If t_1 and t_2 are terms, then $(t_1 = t_2)$ is a formula
   • $s[t_1 = t_2] = \{i \in s \mid i(t_1) = i(t_2)\}$

20. In section 4.5 we shall see that we need one more basic type of expression: some kind of demonstratives.

21. If the language would contain expression with presuppositions, they would be another source of partiality of updates. See the references given on page 142.

In case one of the terms in an atomic formula is a variable that is not in the set v in the referent system of s, i.e., the set of variables that have been introduced in s, the interpretation of that variable does not exist, and hence the update of s with that formula does not exist either. Atomic updates are partial, unless they contain no variables.

From their definition it is clear that atomic updates are distributive. We look at each $i \in s$ to see whether in w the interpretation of the predicate holds of that of the term(s). If so, i will be in the output state, if not, i is eliminated. So, atomic updates are also purely eliminative, and hence classical.

The effect of updating a state s, e.g., with the formula Pa, is that all possibilities $i \in s$ will be eliminated that contain a world w in which the denotation of a is not in the set of objects denoted by P. If there is at least one such possibility i in s, i.e., if it is compatible with our information that a has the property P, then s allows Pa. (And there are certainly states where this is so, which means that Pa is consistent.) Updating such a state s with Pa results in a state $s[Pa]$ which supports Pa. Then we have the information that a has the property P. This shows the rather trivial fact that $Pa \models Pa$. (Still, it is worth mentioning, because we will see that not every formula entails itself.)

Let us also look at an example of an identity statement. Unlike the previous example, the formula $x = a$ is a partial update. For the existence of the update of s with $x = a$ it is presupposed that the variable x has been introduced already. If s meets this constraint, then $s[x = a]$ is the result of eliminating those possibilities $i \in s$ in which the value assigned to x is not the object that *in that possibility* is denoted by a. In other words, in all possibilities i that remain in $s[x = a]$, the value of x is set to $w(a)$, where w is the world of i. In $s[x = a]$ we have the information that x is a. But notice that it is still possible that in different possibilities i and i' with different worlds w and w', x is assigned a different value, because it need not be the case that $w(a) = w'(a)$ for all w and w'.

This means that identity statements such as $a = b$ are (epistemically) *contingent*. They can provide real information, i.e., they can constitute a real update. This is not at odds with the common view that proper names are rigid designators. The property of rigidness is a metaphysical, not an epistemic one. Assuming rigid designation, identity statements such as $a = b$, if true, are necessarily true, but the necessity involved is metaphysical necessity, not epistemic necessity.[22]

The above discussion gives rise to the following general characterization of the properties of atomic updates:

Fact 10
Atomic updates are consistent and coherent partial classical updates

The fact that atomic updates are classical, means that the interpretation of atomic formulae is not inherently dynamic.

4.1.3 (Re-)Assignment

In section 3.12, we have added variables to our referent systems, and have given an informal description of what the effects are of using an existential quantifier. As we have seen, this involves: adding a variable to the set of variables in the referent system; at the same time adding the next peg to the set of pegs; associating the variable with the new peg; assigning an object to the peg.

The operation that establishes this has two parameters: a variable, x, and an object, d. It operates on possibilities, and in an extended sense on states. The definition is as follows.

Definition 30
Let $i = \langle v, n, r, g, w \rangle \in I, s \in S, x \in Var, d \in D$.
1. $i[x/d] = \langle v \cup \{x\}, n+1, r[x/n], g[n/d], w \rangle$
2. $s[x/d] = \{i[x/d] \mid i \in s\}$

Concerning the first clause we note the following. The new association function $r[x/n]$ is like r, except that x is (re-)associated with the new peg n. This means that $r[x/n]$ is a real extension of r. And similarly for $g[x/d]$, which is a real extension of g. Furthermore, $v \cup \{x\}$ is an extension of v, a real one if x was not already in v. Finally, $n < n+1$. Since the world w in $i[x/d]$ is the same as in i, it follows that $i[x/d]$ will always be a real extension of i.

The second clause tells us that (re-)assignment in an information

22. The worlds in an information state are epistemic or doxastic alternatives only. The present paper is only about the corresponding notions of modality, and not about ontological or metaphysical modality. We are well aware of the fact that the question of how to model both kinds of modality within one and the same framework remains, but our aim is more modest. It is to model what is involved in learning to identify the individual a name refers to.

state is done pointwise with respect to its possibilities. To each $i \in s$ corresponds a unique $i[x/d]$ in $s[x/d]$. Since we saw already that each such $i[x/d]$ is a real extension of i, we know that $s[x/d]$ is an extension of s. It will be a real extension, except when s is the absurd state, because $\emptyset[x/d] = \emptyset$.

At the same time, all information about the world and the discourse information of s is retained in $s[x/d]$. Only a new discourse question about the possible value of the newly introduced peg, and a particular 'hypothetical' answer to that question, that it is the object d, have been added. This means that s subsists in $s[x/d]$.

Thus the following four facts hold for (re-)assignment.

Fact 11

1. $i < i[x/d]$
2. $s \leq s[x/d]$
3. If $s \neq \emptyset$, then $s < s[x/d]$
4. $s \sqsubseteq s[x/d]$

With regard to the composition of g and r, which gives us the objects in the domain that are the values of the variables *via* their association with the pegs, it is important to note that it need not hold that for all $x \in v\colon g(r(x)) = g'(r'(x))$. In particular this need not hold if the variable x that is being introduced was already present in v. This means that if we allow re-introduction of variables and would assign them values *directly*, there would be no guarantee that re-assignment would result in a state that is an extension of the original one. But then we would not be able to define a suitable notion of extension between possibilities. And since that notion is the basis of all other notions we need, we would be left empty-handed. Using pegs, however, we are able to define a proper notion of extension, also for re-assignment. This illustrates once more the technical importance of the pegs, which is goes hand in hand with a certain intuitive appeal, and philosophical and cognitive significance. Unfortunately, we cannot dwell upon this issue any further here.

We end this section with the following observation. (Re-)assignment is a function from information states to information states, it is an update. We observe that given its definition, and the facts listed above, it holds that:

Fact 12

A (re-)assignment constitutes a distributive and non-eliminative update, that is allowed and supported by every non-absurd information state.

The fact that (re-)assignment is non-eliminative is important, it means that it is not classical, and that hence it involves a genuine dynamic element.

4.1.4 Existential quantification

By now we have set the stage for the interpretation of existential quantification. We start with the definition.

Definition 31

- If ϕ is a formula, and x a variable, then $\exists x\phi$ is a formula
- $s[\exists x\phi] = \cup_{d\in D}(s[x/d][\phi])$

If we update a state s with $\exists x\phi$, we pick an object d from the domain, and we (re-)assign d to x in s. The state $s[x/d]$ that results from this is updated with ϕ. After we have done this for every object d, we collect the results.

By way of an example, let us update a state s with $\exists xPx$. The result will be a state s' which has the following features. First of all, its referent system will contain a new peg, which is associated with x, which will be in the set of variables of the new referent system. Note that this new referent system is uniform in the new state: it is the same in every possibility in s'.

Second, if in the world w of some possibility $i \in s$ no object has the property P, i.e., if $i(P) = w(P) = \emptyset$, then no extensions of i will appear in s'. If in all possibilities $i \in s$ there is at least one object in its world that has the property P, then s already supports $\exists xPx$. And if in no possibility $i \in s$ there is an object that has the property P, then s forbids $\exists xPx$. Otherwise s simply allows $\exists xPx$ and in s' some information about the world has been gained, viz., that there is some object which has the property P, which is what one would expect. For all worlds in which no object has the property P will have been eliminated as possible alternatives of what the world might be like.

Third, for each possibility $i \in s$, and for each object d that in i has the property P, a real extension $i[x/d]$ of i will appear in s'. So for each possibility i in s there will appear as many extensions $i[x/d]$ of i in $s[\exists xPx]$ as there are objects with the property P in the world of i. This means that if $\exists xPx$ is allowed at all in s, it cannot fail to produce a new state that is a real extension of s. So, even though Px itself is an eliminative update, $\exists xPx$ is not. The possibilities in s' are not a subset of the possibilities in s, but real extensions of a subset of the possibilities in s. This illustrates the non-eliminative nature of existential quantification, that gives it its dynamic impact.

Note that this holds even if s supports $\exists xPx$ already. What guarantees this is the introduction of a new peg. To see this, suppose we update a state s which allows it twice in a row with $\exists xPx$. After the first update, we get into $s[\exists xPx]$, which supports $\exists xPx$. Even then, the second update leads to a real extension, but only because there is a new peg introduced. That nothing essential has happened after the second update is reflected by the fact that the states that results after the first and after the second update, though not the same, are highly similar:

$$s[\exists xPx] \neq s[\exists xPx][\exists xPx]$$

$$s[\exists xPx] \approx s[\exists xPx][\exists xPx]$$

Things are different if we compare $\exists xPx$ and $\exists yPy$. Here, too, we do have that once we have updated with the one, the other cannot fail to be supported, and vice versa. However, the two are not equivalent:

$$\exists xPx \models \exists yPy$$

$$\exists xPx \not\equiv \exists yPy$$

These two formulae can not be substituted for each other, since they provide different discourse information. The one introduces another variable than the other, and that makes a difference. For example, updating with $\exists xPx$ guarantees that an update with Px exists. But updating with $\exists yPy$ does not.

Actually, updating a state s with $\exists xPx$, not just guarantees that an update with Px exists, it actually leads to a state that supports it:

$$\exists xPx \models Px$$

This is a simple, but significant instance of dynamic entailment. A quantifier in a premiss can bind a 'free' variable in the conclusion. A static notion of entailment can do no such tricks. Notice that the entailment also holds the other way around:

$$Px \models \exists xPx$$

Yet the two are not equivalent.

Although the notion of equivalence is sensitive to which quantifiers have been used, it does not care about the order in which they have been introduced:

$$\exists x\exists yRxy \equiv \exists y\exists xRxy$$

Finally, we note that even formulae which are universally supported need not be equivalent:

$$\models \exists x(x = x)$$
$$\models \exists x \exists y(x = y)$$
$$\exists x(x = x) \not\equiv \exists x \exists y(x = y)$$

All this pertains to the non-eliminative nature of existential quantification, which it inherits from (re-)assignment.

We noticed already that our example $\exists x Px$ is a real update (in every non-absurd state). But it can be shown that this holds in general.

Fact 13
If ϕ is an update, then for all s: if $s \neq \emptyset$ and $s[\phi]$ exists, then $s < s[\exists x\phi]$

Suppose ϕ is an update, i.e., for all s: $s \leq s[\phi]$. From fact 11 we know that for all s: if $s \neq \emptyset$, then $s < s[x/d]$. We arrive at $s[x/d][\phi]$ by applying the composition of the updates constituted by the (re-)assignment and by ϕ. The composition of a real update and an update cannot fail to be a real update in a non-absurd state. This means that $\forall s$: if $s \neq \emptyset$ then $s < s[x/d][\phi]$. The interpretation of $\exists x\phi$ amounts to taking unions of the results of $s[x/d][\phi]$ for different d. In general it holds that If $s < s'$ and $s < s''$ then $s < s' \cup s''$, if $s' \cup s'' \in S$. Hence, $\forall s$: if $s \neq \emptyset$ then $s < \exists x\phi$, if ϕ is an update.[23]

As for distributivity the following holds:

Fact 14
If ϕ is a distributive update, then $\exists x\phi$ is a distributive update.

We have seen in fact 12 that (re-) assignment constitutes a distributive update. The composition of two distributive updates is also distributive. And taking the union of the outcomes of $s[x/d][\phi]$ for different $d \in D$ cannot distort this picture.

---

23. We note in passing that it is not guaranteed that if s and s' are information states, then $s \cup s'$ is also an information state. Of course, it will hold that $s \cup s' \subseteq I$, but it is not guaranteed that all $i \in s \cup s'$ have the same referent system. If s and s' do not have the same referent system, then not all possibilities in their union will have the same referent system, and hence, $s \cup s'$ is not an information state. In the union we take in the case of existential quantification, this situation cannot arise. All the alternatives we look at do have the same referent system. Note that $s \cap s'$ always happens to constitute an information state. For if s and s' have different referent systems, their intersection will be empty, and $\emptyset$ is an information state, be it not the most interesting one.

We will see in section 4.3 that if ϕ is non-distributive, then neither is $\exists x\phi$, which will turn out to be rather important.

So, we can conclude the following:

Fact 15
Existential quantification is a non-eliminative real update, it is distributive if its complement is. It is allowed (forbidden) in a state if that state extended with a new peg, under some (every) assignment of an object to that peg, allows (forbids) the complement of the quantifier. It is supported by a state if that state extended with a new peg, under every assignment of an object to that peg, supports the complement.

That existential quantification is not a classical update means that it is inherently dynamic.

4.1.5 Sequencing

One effect of updating a state s with $\exists x\phi$ is that each of the possibilities in the resulting state presents us with a possible value for x, given a certain possibility for the values of the other variables that have been used, and given one way in which the world might be. In short, all information we now have about what might be the alternative values for x is present in the information state that is the output of $s[\exists x\phi]$. If we continu to update with a formula ψ which contains a free occurence of x, this information is still there. This means that the existential quantifier is still operative: it binds such occurrences of variables in ψ.

We write a sequence of updates as a conjunction.

Definition 32

- If ϕ and ψ are formulae, then $(\phi \wedge \psi)$ is a formula.
- $s[\phi \wedge \psi] = s[\phi][\psi]$

Conjunction is a pretty dull operation. It just passes on information. It hands over the output of the first conjunct as input for the second. Still, it has some exciting consequences, such as:

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

An existential quantifier can bind occurrences of variables in formulae that follow it in a sequence, and which hence are outside its scope, with the same force as if they were inside its scope. This is the extended binding power that a dynamic semantics lends to the existential quantifier.

However, note that it does not hold in general that $\exists x\phi \wedge \psi$ and $\exists x(\phi \wedge \psi)$ are equivalent. This is not because there is no extended binding, but because in certain (modal) contexts a difference in meaning results if a variable is bound outside the scope of the quantifier. The following restricted version does hold:

Fact 16
If ϕ and ψ are distributive updates, then $\exists x\phi \wedge \psi \equiv \exists x(\phi \wedge \psi)$

We can not yet give examples that show that the general extended binding equivalence fails, using non-distributive updates, because all updates we have discussed so far are distributive. Sequencing inherits certain properties of its parts. If both conjuncts are distributive (eliminative) then the conjunction is, too. If one of them lacks the property, the conjunction may lack it, too. Since atoms are distributive, and existential quantification preserves distributivity, whatever formula we can construct sofar constitutes a distributive update. We will return to the issue in section 4.3.

Another thing to be noticed about sequencing, is that it is associative. Given that the interpretation of a sequence of sentences, is the composition of the updates associated with the sentences themselves, this follows immediately from the fact that updates are functions and that function composition is associative.

Fact 17
$((\phi \wedge \psi) \wedge \chi) \equiv (\phi \wedge (\psi \wedge \chi))$

This equivalence allows us to leave out brackets in a series of conjunctions. For example, instead of $((\phi \wedge \psi) \wedge \chi)$ we may write $(\phi \wedge \psi \wedge \chi)$.

Function composition is not generally commutative, and update sequencing is not generally commutative either. Here we *can* come up with a counterexample, which at the same time is a counterexample against idempotency of entailment.

$$x = a \wedge \exists x(x = b) \not\equiv \exists x(x = b) \wedge x = a$$

$$x = a \wedge \exists x(x = b) \not\models x = a \wedge \exists x(x = b)$$

The last of these two facts is perhaps the more obvious one. There are certainly possible information states in which the information is present that

a and b refer to different objects, and which contain at least one possibility in which the value assigned to x is the denotation of a in that possibility. Such a state will allow $x = a \wedge \exists x(x = b)$, but the state that results from updating with it, will in fact forbid the conclusion that $x = a \wedge \exists x(x = b)$.

The reason why idempotency fails is that the quantifier in the premiss binds the free variable in the conclusion. Similarly, because the variable x occurs free in the first conjunct of $x = a \wedge \exists x(x = b)$, and since the quantifier $\exists x$ is actively present in the second, commuting the conjuncts does not yield an equivalent result. What once was a free variable, is suddenly a bound one. (If we read the equivalence statement from left to right!) It is precisely the extended binding power of the existential quantifier that prevents commutativity to hold. In a state where we know a and b to denote different objects $x = a \wedge \exists x(x = b)$ can be allowed (depending on the possible values of x), whereas $\exists x(x = b) \wedge x = a$ will be forbidden. Or, to take a much simpler case, there are certainly states in which the variable x has not been introduced yet. Such a state might very well allow $\exists x(x = b) \wedge x = a$, but would not be able to allow $x = a \wedge \exists x(x = b)$, because the update with it does not exist.

Of course, in many cases the order of a sequence of updates does not matter. Sequences of classical updates are a case in point.

Fact 18
If ϕ and ψ are classical updates, then $\phi \wedge \psi \equiv \psi \wedge \phi$

Notice that this fact states a sufficient, but not a necessary condition for commutativity. We will not bother to state the precise conditions under which commutativity holds.[24] Rougly speaking, and restricting ourselves to distributive updates, if there is no overlap in the variables bound by ϕ and the free variables occurring in ψ, and vice versa, then we can commute $\phi \wedge \psi$.

We conclude again with a general characterization:

Fact 19
Sequencing inherits its properties from its conjuncts. It is eliminative (distributive), if both of its conjuncts are. A conjunction is allowed (forbidden) (supported) by an information state, if the first conjucnt is allowed (forbidden) (supported) by that information state, and if the second conjunct is allowed (forbidden) (supported) by the information state that results from

24. See Groenendijk and Stokhof 1991 for extensive discussion.

updating the orginal state with the first conjunct.

Sequencing is, one might say, 'passively' dynamic. It does nothing to stop the dynamics of what it operates on, but it does not introduce dynamic effects of its own either.

4.1.6 Negation

With negation it is the other way around: it is 'actively' static, i.e., it blocks the dynamics of whatever is in its scope.

Definition 33

- If ϕ is a formula, then $\neg\phi$ is a formula.
- $s[\neg\phi] = \{i \in s \mid \neg\exists i' \colon i \leq i' \text{ and } i' \in s[\phi]\}$

A negation $\neg\phi$ eliminates those possibilities in s of which some extension can be found in the state that would result after updating s with ϕ. This means that whatever extensions quantifiers inside the scope of a negation may bring about, these are nullified. Negation is a purely eliminative update, no matter what the nature of it complement is.

If the complement of negation is itself eliminative, then $s[\neg\phi]$ boils down to $\{i \in s \mid i \notin s[\phi]\}$. And if ϕ is classical, things may get even simpler, as can be seen from the negation of an identity statement:

Definition 34

- $t_1 \neq t_2 =_{\mathrm{df}} \neg(t_1 = t_2)$
- $s[t_1 \neq t_2] = \{i \in s \mid i(t_1) \neq i(t_2)\}$

Consider an example: $\neg(\exists x Px \wedge Qx)$. Updating a state s with this formula eliminates those possibilties $i \in s$ such that there is some extension i' of i that would be in the result of updating s with $\exists x Px \wedge Qx$. That is to say, those possibilities in s are eliminated that would allow there to be something that has both the property P and the property Q. And that is how it should be.

Of course, in establishing $s[\exists x Px \wedge Qx]$, the extended binding of the quantifier is operative. That is why we have to look at *extensions* $i' \in s[\exists x Px \wedge Qx]$ of possibilities i in the initial state s. But the net result is a state consisting of certain i in the original s. This means that, in the end, no changes in the referent system are made. The new peg that the quantifier introduces in calculating $s[\exists x Px \wedge Qx]$, does not appear in the end result $s[\neg(\exists x Px \wedge Qx)]$.

In section 2 we saw that this feature of negation is backed up by linguistic facts.[25] It also makes sense from a logical point of view. The formula $\neg(\exists x Px \wedge Qx)$ says that there are no objects that have both the property P and the property Q. So, what possible values of (the peg associated with) x are we supposed to keep track of?[26]

One fact to note about negation is that a characteristic property of classical negation, the law of double negation, does not hold in general, but only in a restrictive version:

Fact 20
If ϕ is an eliminative update, then $\neg\neg\phi \equiv \phi$

Counterexamples with non-eliminative updates are easily provided:

$$\exists x Px \not\equiv \neg\neg\exists x Px$$

It is the difference in discourse information that they carry, that makes these two formulae non-equivalent: whereas the quantifier in $\exists x Px$ can still bind variables in sentences to come, the one $\neg\neg\exists x Px$ can not.

It does not hold in general that ϕ and $\neg\neg\phi$ mutually entail each other:

Fact 21
1. $\neg\neg\phi \models \phi$
2. If ϕ is an eliminative update then $\phi \models \neg\neg\phi$

A counterexample with a non eliminative update is provided by:

$$Px \wedge \exists x\neg Px \not\models \neg\neg(Px \wedge \exists x\neg Px)$$

Again, the reason is that the quantifier in the premise binds the free variable in the conclusion.

Although perhaps akward from a (classical) logical point of view, from the perspective of natural language semantics and reasoning this is a

25. Which is not to say that it holds unconditionally. See Groenendijk and Stokhof 1991, p. 89–92, Dekker 1993b, ch. 2 for some discussion.
26. Answer: we might preserve the ones that got removed. But these we can not store in the new state just like that. One way or another, we should store them as values we used to have, but which are now eliminated. As the states are set up right now, there is no room for this type of information, but maybe one should make room for it.

nice feature of dynamic entailment. In argumentative discourse, anaphoric relations between binding expressions in the premisses of an argument, and variable-type expressions (such as pronouns) in a conclusion, occur all the time. Trying to do without them can become rather cumbersome. Dynamic entailment makes for an easier and more natural formalization of such discourses.

All this has to do with the eliminative nature of negation. What about distributivity? Again, we can not provide any counterexamples yet, because anything that negation can operate on sofar is distributive. Given the way it is defined, the negation of a non-distributive update is non-distributive itself.

Let us sum up:

Fact 22
Negation is an eliminative update, it is distributive when its complement is. It is allowed when its complement is not supported, it is forbidden if its complement is supported, and it is supported if its complement is forbidden.

Negation turns everything upside down.

4.1.7 Implication, disjunction and universal quantification

Having defined the updates constituted by the atoms, existential quantification, sequencing and negation, we have defined the predicate logical fragment of the language we are after. The one new thing that remains is the *might*-operator, which we will have a look at in the next section.

Of course there are the usual other operations of predicate logic, viz., implication, disjunction and universal quantification, but we can introduce them by defining them in terms of the operations that we already have. The definitions are the usual ones.[27]

Definition 35
1. $(\phi \rightarrow \psi) =_{\mathrm{df}} \neg(\phi \wedge \neg\psi)$
2. $(\phi \vee \psi) =_{\mathrm{df}} \neg(\neg\phi \wedge \neg\psi)$
3. $\forall x\phi =_{\mathrm{df}} \neg\exists x\neg\phi$

27. We leave material equivalence $\leftrightarrow$ out of consideration here. We suspect that introducing it would mean introducing it as a separate basic operation. Defining it in the ususal way, i.e., as a conjunction of mutual implications, would not work. It would result in a connective that would allow for binding from right to left. We see no easy way around this.

If we calculate their interpretations given these definitions we arrive at the following:

Fact 23

1. $s[\phi \to \psi] = \{i \in s \mid \forall i'$: if $i \leq i'$ and $i' \in s[\phi]$, then $\exists i'' : i' \leq i''$ and $i'' \in s[\phi][\psi]\}$
2. $s[\phi \vee \psi] = \{i \in s \mid \exists i' : i \leq i'$ and $i' \in s[\phi]$ or $i' \in s[\neg\phi][\psi]$
3. $s[\forall x\phi] = \{i \in s \mid \forall d \in D : \exists i'$ such that $i \leq i'$ and $i' \in s[x/d][\phi]\}$

Implication amounts to the following. If we update a state with $\phi \to \psi$, then a possibility will remain in case if it subsist after an update with ϕ, then all its survivors after updating with ϕ subsist after a further update with ψ. Disjunction says that if we update a state with $\phi \vee \psi$, those possibilities remain which subsists after an update with ϕ, or after an update with ψ. Universal quantification, finally, says that if we update a state with $\forall x\phi$, then those possibilities remain, which after every (re-)assignment of x to some object, subsists after an update with ϕ.

The first thing to note is that all three of these connectives constitute eliminative updates. This feature they inherit from negation, which in their respective definitions is the outermost operator. This means that quantifiers inside their scope cannot bind free variables in further sentences. And with respect to universal quantification we may note that, given its eliminative nature, it can not bind variables outside its scope itself either. As is the case with negation, to a large extent this feature is borne out by the facts. However, there are exceptions one should be able to deal with properly. But we will not attempt to do so here.[28]

Another observation is that there is no alternative choice for the basic operations. This follows immediately from the eliminative nature of the alternatives. For example, if we were to take implication instead of sequencing as a basic operation, we could not arrive at a proper interpretation for sequencing. If we would define $(\phi \wedge \psi)$ as $\neg(\phi \to \neg\psi)$, we would block the binding of quantifiers inside its scope in further sentences.

One special feature of implication, which also follows directly from the way it is defined, is that although quantifiers inside its scope cannot bind variables outside its scope, quantifiers in its antecedent can bind variables in its consequent. This gives our semantics the possibility to give straightforward translations of the so-called donkey-sentences. The following equiv-

28. See the references cited in footnote 11.

alence presents the essentials of that fact:

$$\exists x Px \to Qx \equiv \forall x(Px \to Qx)$$

However, it does not hold in general that $\exists x\phi \to \psi$ and $\forall x[\phi \to \psi]$ are equivalent. It only holds if ψ is a distributive update.

Fact 24
If ψ is a distributive update, then $\exists x\phi \to \psi \equiv \forall x[\phi \to \psi]$

Again, since all updates defined sofar are distributive, counterexamples to the more general claim have to wait until we have introduced non-distributive updates.

One final thing we may note is that there is a deduction theorem:

Fact 25
$\models \phi \to \psi$ iff $\phi \models \psi$

4.2 Modality

The last basic logical operation that remains to be introduced is the *might*-operator.

Definition 36

- If ϕ is a formula, then $\Diamond\phi$ is a formula
- $s[\Diamond\phi] = \{i \in s \mid s[\phi] \neq \emptyset\}$

Provided that $s[\phi]$ exists, there are only two possible outcomes of updating a state s with $\Diamond\phi$. We either remain in s, or we end up in the absurd state. The first happens if updating s with ϕ does not lead to the absurd state, the second if it does.

In other words, when updating s with $\Diamond\phi$ we remain in s if s allows ϕ; and we reject $\Diamond\phi$ if s forbids ϕ. Every state either supports or forbids $\Diamond\phi$. Having only these two possibilities, $\Diamond\phi$ is a test, and what it tests for is whether ϕ is consistent with our information. And like all tests, $\Diamond\phi$ is eliminative, which means that existential quantifiers inside its scope cannot bind variables outside.

We have been waiting a long time for this. Finally the language is no longer completely distributive: $\Diamond\phi$ is a non-distributive update. Consistency testing essentially involves looking at an information state globally, and not pointwise with respect to the possibilities it contains. We illustrate this with an example.

Let p be a proposition, i.e., $p \in R^0$. Its interpretation in a world w is $w(p)$, where $w(p) \in \{0,1\}$.[29] Application of the clause for atomic formulae gives: $s[p] = \{i \in s \mid i(p) = 1\}$, where $i(p) = 1$ iff $w(p) = 1$ for the world w in i.

Let $s = \{i, j\}, i(p) = 1, j(p) = 0$. Then $s[\Diamond p] = s$, since $s[p] = \{i\}$, and hence $s[p] \neq \emptyset$. On the other hand, $\{i\}[\Diamond p] = \{i\}$ and $\{j\}[\Diamond p] = \emptyset$. So, performing the update expressed by $\Diamond p$ pointwise on the elements i and j of s, and taking the union of the two results, gives $\{i\}$. And since $\{i\} \neq s$, we see that we do not get the same outcome as when we apply $\Diamond p$ to s globally. This proves that $\Diamond\phi$ is not a distributive update.

Representing *might* as $\Diamond$, as we do here, interprets it as a consistency test. But this is only one aspect of its meaning. If you are told 'It might be raining today', this may very well constitute *real* information, on the basis of which you could decide, for example, to take an umbrella when going out. This additional aspect of the meaning of *might* is not dealt with in the present semantics, which merely takes into account that if you hear that it might rain, you check whether it is possible with respect to the information you have that it rains. If so, nothing happens. But, apparently, that is not all there is to it.

On the other hand, what our semantics does account for is the following observation. Suppose again that you hear that it might be raining. And suppose furthermore that the information you have tells that this is not so ('No! Look outside! The sun is shining!'), i.e., suppose that the information of it raining outside is inconsistent with the information you have. Then, in all likelihood, you will not accept the remark just like that, you will start arguing. It is this aspect of the meaning of *might* that is accounted for by the semantics we have given above. If ϕ is inconsistent with s, updating with $\Diamond\phi$ results in $\emptyset$. And precisely because you do *not* want to end up in the absurd state, you do *not* just update your information state, but will start arguing with whoever tries to tell you that $\Diamond\phi$. It is for this reason that we say that if $s[\phi] = \emptyset$, s *forbids* ϕ.

On this score, the semantics for *might* does well. It predicts correctly that if you are told that it might be raining, when you can look out and see

29. Since p is a 0-place predicate, $w(p) \subseteq D^0$. D^0 is the set of 0-tuples, which is $\{\langle\rangle\}$. The empty sequence $\langle\rangle = \emptyset$. So, $D^0 = \{\emptyset\}$. This set has two subsets: itself, $\{\emptyset\}$, and $\emptyset$. And these are the set theoretical definitions of 1 and 0, respectively. From this one can figure out that the next two sentences in the text are true also.

that the sun is shining, you will start arguing about it. And this, in turn, means that consistency testing is involved in the semantics of *might*.

That *might*-ϕ may in some states provide a real update of information is not accounted for as yet. Let us give two examples where this might be the case.

One point at which this effect may occur is in the information that the hearer has 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 allows p. Since this type of higher-order information is left out of consideration here, this kind of update effect is not accounted for.

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

(17) It might rain. It would ruin your blue suede shoes.[30]

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 you are wearing will be ruined, which could be a real update, and not just a test.

This phenomenon, which is known under the name of *modal subordination*, is a central feature of the meaning of natural language modalities.[31] (And, actually, it can be used as a first rate argument in favor of a dynamic treatment of it.) Nevertheless, we will (almost) completely ignore it in the present paper. For the moment it suffices to indicate that there is more to *might* than the semantics presented in this paper covers. On the other hand, we also hope to have indicated that consistency testing is indeed an essential ingredient of its meaning.[32]

Let us illustrate the role of *might* in discourse somewhat more by looking at some examples of sequences in which this modality figures.

In section 2 we discussed the following two sequences of sentences:

(18) a. It might be raining outside [...] It isn't raining outside.

---

30. Just like stepping on them would.

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

32. Another aspect of the meaning of *might* that is not accounted for is that a sentence of the form $\Diamond p$ may also have the effect that one comes to realize the possibility of p as such. The eliminative nature of the present framework prohibits an account of this: eliminativity attributes to language users 'omniscience' with respect to what the possibilities are.

b. It isn't raining outside [...] *It might be raining outside.

The example shows that order matters in updating an information state with a sequence of sentences. In our logical language, (18a) corresponds to $\Diamond p \wedge \ldots \wedge \neg p$, and (18b) to $\neg p \wedge \ldots \wedge \Diamond p$. For the moment we ignore what might occur on the dots between the two sentences in the sequences, but we return to this question below. We focus on the two formulae $\Diamond p \wedge \neg p$ and $\neg p \wedge \Diamond p$. In ordinary static modal logics, these two formulae are equivalent, and hence these systems can not account for the difference between (18a) and (18b). In our dynamic semantics, the difference between the two is accounted for by the following fact:

Fact 26

1. $\Diamond \phi \wedge \neg \phi$ is consistent
2. $\neg \phi \wedge \Diamond \phi$ is inconsistent

First we prove the acceptability of $\Diamond \phi \wedge \neg \phi$. To succeed it is enough to come up with some formula ϕ and some state s such that s allows $\Diamond \phi \wedge \neg \phi$. To make things easier, we choose a proposition p for ϕ. So we have to come up with some state s which allows $\Diamond p \wedge \neg p$.

Let $s = \{i_1, i_2\}$, where $i_1(p) = 1$ and $i_2(p) = 0$. We note at the outset that this means that i_1 and i_2 should contain two different worlds w_1 and w_2 such that $w_1(p) = 1$ and $w_2(p) = 0$. By the definition of the extension relation, this means that $i_1 \not\leq i_2$ (nor the other way around).

By the definition of sequencing: $s[\Diamond p \wedge \neg p] = s[\Diamond p][\neg p]$. By the definition of *might*: $s[\Diamond p] = s$, since $s[p] = \{i_1\} \neq \emptyset$. So, $s[\Diamond p][\neg p] = s[\neg p]$. According to the definition of negation: $s[\neg p] = \{i \in s \mid \neg \exists i' \in s[p]\colon i \leq i'\}$. According to the atomic clause: $s[p] = \{i_1\}$.

For i_1 it holds that $\exists i' \in s[p]\colon i_1 \leq i$, viz., i_1 itself. So, $i_1 \notin s[\neg p]$. For i_2 it does *not* hold that $\exists i' \in s[p]\colon i_2 \leq i'$. Only $i_1 \in s[p]$ and $i_2 \not\leq i_1$, as we noted at the outset. Hence, $i_2 \in s[\neg p]$. So, since $i_1 \notin s[\neg p]$, and $i_2 \in s[\neg p]$, and since according to the definition of negation for $i \in s[\neg p]$, it should hold that $i \in s$, we get that $s[\neg p] = \{i_2\}$.

Hence, since we saw already that

$$s[\neg p] = s[\Diamond p][\neg p]$$

and that

$$s[\Diamond p][\neg p] = s[\Diamond p \wedge \neg p]$$

we have shown that

$$s[\Diamond p \wedge \neg p] = \{i_2\} \neq \emptyset$$

This means that s allows $\Diamond p \wedge \neg p$. And if there is some such s, then $s \in S$, $\Diamond p \wedge \neg p$ is consistent. ■

Next we prove the inconsistency fact by showing that the assumption that $\neg\phi \wedge \Diamond\phi$ is inconsistent, leads to a contradiction.

We may assume that ϕ is an update, which means that $\forall s\colon s \leq s[\phi]$ Which means that $\forall s\colon \forall i \in s\colon \exists i'\colon i \leq i'$ and $i' \in s[\phi]$. If $s[\phi] \neq \emptyset$ this means that $\exists i\colon i \in s$ and $\exists i'\colon i \leq i'$ and $i' \in s[\phi]$.

Suppose $\neg\phi \wedge \Diamond\phi$ is consistent. By the definition of acceptability this means that there is some state s such that s allows $\neg\phi \wedge \Diamond\phi$.

By the definition of acceptance this requires that $s[\neg\phi \wedge \Diamond\phi] \neq \emptyset$. By the definition of sequencing this means that $s[\neg\phi][\Diamond\phi] \neq \emptyset$. By the definition of *might* this means that $s[\neg\phi][\phi] \neq \emptyset$. By the definition of negation this means that $\{i \in s \mid \neg\exists i'\colon i \leq i'$ and $i' \in s[\phi]\}[\phi] \neq \emptyset$.

As we saw above, using the assumption that ϕ is an update, this would require $\exists i\colon (\neg\exists i'\colon i \leq i'$ and $i' \in s[\phi])$ and $(\exists i'\colon i \leq i'$ and $i' \in s[\phi])$. But this is a contradiction. Since the assumption that $\neg\phi \wedge \Diamond\phi$ is consistent leads to a contradiction, we have shown that $\neg\phi \wedge \Diamond\phi$ is inconsistent. ■

That $\Diamond\phi \wedge \neg\phi$ is consistent whereas $\neg\phi \wedge \Diamond\phi$ is not, means that the two are not logically equivalent. In the previous section we already gave a counterexample against general commutativity with non-eliminative updates, here we have a counterexample with non-distributive updates.

Notice that $\Diamond p \wedge \neg p$ provides another counterexample against idempotency of entailment:

$$\Diamond p \wedge \neg p \not\models \Diamond p \wedge \neg p$$

And this time, this has not to do with non-eliminativity, but with non-distributivity.

We promised to return to the dots in such sequences as in (18). The dots are used to indicate that some time elapses between the utterances of the two sentences. In the meantime information from other sources may become available. For example, someone might open the blinds and look outside. To see why this is important, suppose we leave out the dots, and assume that the two sentences are uttered in immediate sequence, by the same speaker, so that virtually nothing can happen in between. In such a case we would be inclined to judge that the first sequence can not be consistent either.

To see why, we dramatize things a bit by making the sequence into a real conjunction:[33]

(19) *It might be raining outside and it is not raining outside.

Obviously this can not be a correct utterance. But we just saw that its logical translation is consistent. How can that be?

The answer is the following. First of all, acceptability is a hearer oriented notion. Secondly, in modeling the information of the hearer, we do not include any information she might have about the information of the speaker. Given this way of modeling things, saying that $\Diamond p \wedge \neg p$ is consistent means nothing more, or less, than that a hearer can be in an information state in which she can consistently update her information subsequently with $\Diamond p$ and $\neg p$. And from that perspective, $\Diamond p \wedge \neg p$ is indeed a consistent update. (And $\neg p \wedge \Diamond p$, for example, is not.)

This does *not* mean that a *speaker* could utter this sequence correctly. He can not, unless in between uttering the first and the second sentence new information from other sources has become available to him. For a speaker to be able to utter the sequence correctly, his information should support it. This means that there should at least be some information state that supports it. In other words, $\Diamond p \wedge \neg p$ should be coherent, but it is not:

Fact 27
$\Diamond \phi \wedge \neg \phi$ is incoherent

For a state s to support $\Diamond \phi \wedge \neg \phi$, s should support $\Diamond \phi$, and s updated with $\Diamond \phi$, should support $\neg \phi$. For s to support $\Diamond \phi$, s should allow ϕ (unless s is the absurd state). If so, s updated with $\Diamond \phi$ is s again. Then s should also support $\neg \phi$. But to support $\neg \phi$ is to forbid ϕ. No state can manage to allow and at the same time forbid a formula. Hence, $\Diamond \phi \wedge \neg \phi$ is incoherent.

If our hearers would have information about the information of speakers, they would realize this. And recognizing the utterance as necessarily incorrect, they would have a good reason to forbid it.

So we are on the right track. Since we can predict the incoherence of $\Diamond \phi \wedge \neg \phi$, even though it comes out as consistent, we can be rest assured that when we do take information of hearers about the information of speakers into account, things will work out fine.

33. Things are different, of course, if the sequence reads 'It might have been raining, but it isn't', but we are not talking about 'might have been', but about 'might'. About the former we have nothing to say in the present paper.

In the present context, it is important to be aware of this. If a sequence of sentences is declared consistent for a hearer, this does not mean that the sequence is also coherent. To get intuitions right, we should reckon with the possibility that the different sentences in a sequence come from different sources, or... that they come from one and the same source, but that in between uttering them, new information from outside has come up. The importance of the dots should be clear by now. If $\phi \wedge \ldots \wedge \psi$ is to be really consistent, this means that $\phi \ldots \psi$ is consistent, and that ϕ is coherent in some state s, such that ψ is coherent in some state $s'\colon s \leq s'$.

Usually, one defines a $\Box$-operator in terms of the $\Diamond$-operator. Let us not keep behind.

Definition 37

- $\Box\phi =_{\mathrm{df}} \neg\Diamond\neg\phi$
- $s[\Box\phi] = \{i \in s \mid s[\phi] \simeq s\}$

Whereas $\Diamond\phi$ test for acceptance, $\Box\phi$ test for support. If s supports ϕ, then $s[\Box\phi] = s$, if it does not, the absurd state results. So, allowing $\Box\phi$ means supporting ϕ, and not supporting ϕ means forbidding $\Box\phi$. One way to read $\Box\phi$ is as ... *So,* ϕ.

We do not have much to say about this, except that we note that the two basic notions of acceptance and support correspond to the $\Box$ and the $\Diamond$.

Fact 28

Our modal operators are tests. They test for acceptance and support. They are allowed and supported if what they test is so, they are forbidden if what they test is not so.

The modal operators are pretty crude. How cute they are, we will see in the next section, when we see how they co-operate with quantification.

4.3 Coreference and modality

In this section, we turn to the descriptive heart of the matter of the present paper, viz., the interplay between indefinites and modals, i.e., between existential quantification and the $\Diamond$-operator.

It is a remarkable fact that the extended binding power that its dynamic interpretation lends to an existential quantifier allows it also to bind variables which are outside its own scope and inside the scope of a modal operator. For example, the variable in the second conjunct of $\exists x Px \wedge \Diamond Qx$, is bound by the existential quantifier in the first conjunct. Similarly, in the implication $\exists x Px \rightarrow \Box Qx$ the quantifier in the antecedent binds the variable in the consequent.

4.3.1 The case of sequencing

However, whereas in the predicate logical fragment of the language, in which only distributive updates occur, extended binding validates the equivalence of $\exists x\phi \wedge \psi$ and $\exists x(\phi \wedge \psi)$, this does not hold in the full language with epistemic modal operators. In general, $\exists x\phi \wedge \Diamond\psi$ and $\exists x(\phi \wedge \Diamond\psi)$ have a different meaning, i.e., they may take one and the same information state into different states.

In other words, extended binding must be distinguished from the equivalence just mentioned. It may license it, in which case one could say that being bound inside a quantifier's scope, and being bound outside, are equivalent properties of variables. But it may also be the case that there is extended binding, yet the equivalence does not hold, which means that the properties are different.

If what we have said is right, then we should be able to prove the following:

Fact 29
$\exists x\phi \wedge \psi \not\equiv \exists x(\phi \wedge \psi)$

To do so we have to specify instances for ϕ and ψ, and a set of possibilities I, based on a domain D and a set of possible worlds W, such that there is a state s in the set of information states S based on I, such that $s[\exists x\phi \wedge \psi] \not\approx s[\exists x(\phi \wedge \psi)]$.

As for our choice of ϕ and ψ, consider (20a) and (21a), which can be paraphrased as (20b) and (21b):

(20) a. $\exists x(Px \wedge \Diamond Qx)$
b. There is someone who has the property P and might have the property Q.

(21) a. $\exists x Px \wedge \Diamond Qx$
b. There is someone who has the property P. He might have the property Q.

What happens if we update with (20) is the following. The existential quantifier introduces a new peg which is (re-)associated with x. Its possible values have two properties: the property P, and the property of being an object such that our information allows that it has the property Q. The latter implies that if some object lacks the property Q in every world compatible with our information, then after updating with (20), it is not among the possible values of the newly introduced peg associated with x.

In the case of the sequence of sentences (21), things are different. The first conjunct introduces a new peg which is (re-)associated with x. Its possible values have the property P. The next sentence tests whether our information allows that the value of the new peg is an object that has the property Q. This means that among these possible values there should be some object that has property Q in some world compatible with our information. If there is *some* such value and *some* such world, then the test succeeds, and *no* value is eliminated. More in particular, also those objects which in *no* world compatible with our information have the property Q, are *not* eliminated. Precisely in this respect, (21) differs from (20).

Let is set up a situation in which this difference actually comes out. We start rock bottom.

Let our non-logical vocabulary consist of only the one-place predicates P and Q. Let our domain consist of two objects only. Then we can specify D, W, I and S as follows:

1. $D = \{d_1, d_2\}$
2. W = the set of functions w with domain the predicates P and Q, and range the set of subsets of D
3. I is the set of possibilities based on D and W
4. S is the set of information states based on I

Suppose we have the following information. We know that only d_1 has property Q. And we know that only one object has property P, but we do not know which one. The initial state which embodies this information can be specified as follows:

$$s_0 = \{i_1, i_2\} = \{\langle \emptyset, \emptyset, w_1 \rangle, \langle \emptyset, \emptyset, w_2 \rangle\}$$

$$w_1(P) = \{d_1\} \;\&\; w_2(P) = \{d_2\} \;\&\; w_1(Q) = w_2(Q) = \{d_1\}$$

When describing the interpretation of (20) above, we said that if some object lacks the property Q in every world compatible with our information, then after updating with (9), it is not among the possible values of the newly introduced peg associated with x. In our situation this is so, d lacks the property Q in each of our two worlds.

Let us calculate this result. By the definitions of existential quantification and sequencing we find:

$$s_0[\exists x(Px \wedge \Diamond Qx)] =$$

$$s[x/d_1][Px][\Diamond Qx] \cup s[x/d_2][Px][\Diamond Qx] =$$

$$\{i_1[x/d_1], i_2[x/d_1]\}[Px][\Diamond Qx] \cup \{i_1[x/d_2], i_2[x/d_2]\}[Px][\Diamond Qx]$$

We write one of these possibilities out in full:

$$i_1[x/d_1] = \langle\{\langle x, 0\rangle\}, \{\langle 0, d_1\rangle\}, w_1\rangle$$

In this possibility, the value assigned to x via the peg 0, viz., d_1, has the property P, since $d_1 \in w(P)$. In $i_2[x/d_1]$ this is not the case. And, similarly, whereas in $i_2[x/d_2]$ the value assigned to x has the property P, it lacks it in $i_1[x/d_2]$. Returning to our series of equivalences, this means that after applying the definition of the atomic clause to Px two times, we arrive at:

$$= \{i_1[x/d_1]\}[\Diamond Qx] \cup \{i_2[x/d_2]\}[\Diamond Qx]$$

Now we perform the acceptance test $\Diamond Qx$. Only for $\{i_1[x/d_1]\}$ will this test succeed. For d_2 does not have property Q in w_2. Hence we can continue our equivalences as follows:

$$= \{i_1[x/d_1]\} \cup \emptyset = \{i_1[x/d_1]\}$$

Writing this out in full we get:

$$= \{\langle\{\langle x, 0\rangle\}, \{\langle 0, d_1\rangle\}, w_1\rangle\}$$

This means that we have gained information. We now know that d_1 is the individual that has property P, because in no world compatible with our information state does d_2 have the property Q.

Let us now turn to (21), $\exists x Px \wedge \Diamond Qx$. When stating its interpretation above, we claimed that this means that among the possible values of x there should be some object that has property Q in some world compatible with our information. If there is *some* such value and *some* such world, then the test succeeds, and *no* value is eliminated. More in particular, also those objects which in *no* world compatible with our information have the property Q, are *not* eliminated. Precisely in this respect, (21) differs from (20). To show how this difference come about consider again the information state s_0. In this state d_1 has the property Q. Application of sequencing and existential quantification gives the following result:

$$s_0[\exists x Px \wedge \Diamond Qx] =$$
$$(s_0[x/d_1][Px] \cup s_0[x/d_2][Px])[\Diamond Qx] =$$
$$(\{i_1[x/d_1], i_2[x/d_1]\}[Px] \cup \{i_1[x/d_2], i_2[x/d_2]\}[Px])[\Diamond Qx]$$

Calculating Px, we get:

$$= (\{i_1[x/d_1]\} \cup \{i_2[x/d_2]\})[\Diamond Qx] =$$
$$= \{i_1[x/d_1], i_2[x/d_2]\}[\Diamond Qx]$$

Now we perform the acceptance test $\Diamond Qx$. It is succesful: if we update the state with Qx, something remains, viz., $\{i_1[x/d_1]\}$. Hence, the output is:

$$\{i_1[x/d_1], i_2[x/d_2]\}$$

And this differs from the result we obtained for (20). What we have shown is that:

$$s_0[\exists x(Px \wedge \Diamond Qx)] = \{i_1[x/d_1]\}$$
$$s_0[\exists x Px \wedge \Diamond Qx] = \{i_1[x/d_1], i_2[x/d_2]\}$$

Since there is some information state in which (20) and (21) lead to a different outcome, it follows that:

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

■

The difference between (20) and (21) comes out in other ways also. If the discourse $\exists x Px \wedge \Diamond Qx$ is continued with $\Diamond \neg Qx$ nothing happens. The test will succeed, because among the possible values of x there is d_2, which does not have the property Q. Hence, we can also continue our update with $\neg Qx$. In that case we eliminate d_1 as a possible value, and the only possibility that remains is the one in which d_2 is the value of x, and in which the world is w_2, the world in which d_2 has the property P.

In other words, our initial state s_0 both supports and allows $\exists x Px \wedge \Diamond Qx \wedge \Diamond \neg Qx$, and allows, but does not support, $\exists x Px \wedge \Diamond Qx \wedge \neg Qx$.

Notice that it is not a particular feature of the state s_0 that it does not support this sequence, no state does. In other words:

Fact 30
$\exists x Px \wedge \Diamond Qx \wedge \Diamond \neg Qx \wedge \neg Qx$ is consistent, but incoherent

With $\exists x(Px \wedge \Diamond Qx)$ things are different. Both a continuation with $\Diamond\neg Qx$ and one with $\neg Qx$ is forbidden. For updating s_0 with $\exists x(Px \wedge \Diamond Qx)$ left us with d_1 as the only possible value for x. And in no world compatible with our information (for which only w_1 now qualifies) does d_1 have the property Q. So, both $\exists x(Px \wedge \Diamond Qx) \wedge \Diamond\neg Qx$, and $\exists x(Px \wedge \Diamond Qx) \wedge \neg Qx$ are forbidden by our initial state s_0.

This does not mean that $\exists x(Px \wedge \Diamond Qx) \wedge \Diamond\neg Qx$, and $\exists x(Px \wedge \Diamond Qx) \wedge \neg Qx$ are inconsistent per se. It is a characteristic of the situation as we set it up that these formulae are forbidden by it. There are also states in which both are allowed. An example appears if we weaken the information in our sample situation a bit, by removing the information that d_1 has the property Q, but still keep the information that d_2 does not have the property Q. If we do so, two more worlds inhabit our initial information state. One world, w_3, which is like w_1, and another world, w_4, which is like w_2, except that in both Q is true of no individual. In this situation, both $\exists x Px \wedge \Diamond Qx \wedge \Diamond\neg Qx$ and $\exists x Px \wedge \Diamond Qx \wedge \neg Qx$ are allowed. Of course, after updating with $\exists x Px \wedge \Diamond Qx$ it still holds that only d_1 remains as a possible value for x. But now both the world w_1 in which d_1 has Q, and w_3 in which it does not, remain as alternatives. This means that if we perform the test $\Diamond\neg Qx$, it will be successful. Likewise, if we update with $\neg Qx$, the possibility containing w_1, in which d_1 has the property Q is eliminated, but the another possibility, containing w_3, in which x does not have the property Q remains.

4.3.2 Two features

There are two features of the particular semantics of dynamic modal predicate logic presented here which distinguish it from other proposals. Both are essential if we are to obtain proper results for the examples we discussed above, and for several other ones, some of which will be discussed in the next section.

The first feature is that a formula of the form $\Diamond\phi$ is interpreted as a *global* consistency test on information states, not only with respect to information about the world, but also with respect to discourse information.[34] This property of the semantics of $\Diamond$ is imperative in order to get the right

34. In the proposal by Van Eijck and Cepparello (to appear), consistency testing is done distributively with respect to discourse information. Each possible assignment of values to variables is tested seperately for consistency with the information carried by the comple-

results for caes such as discussed in the previous section. Other examples yet to come, will lend further support for this feature.

The second feature of the present semantics that is essential for its empirical adequacy is the way in which existential quantification is dealt with. A rather obvious way of defining definition the effect of updating a state s with $\exists x\phi$ is the following. Start with a global random (re-)assignment of x in s, and next update the state that results with ϕ.

Global (re-)assignment of x in s works as follows. A new peg is introduced. For each object $d \in D$, each possibility in s is extended with d as the value of the new peg, and hence of the variable x associated with it. This amounts to taking $\cup_{d\in D} s[x/d]$.

In terms of that the intuitive definition of existential quantification amounts to this:

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

The definition given in the present paper differs from this one in a subtle, but crucial way, viz., in the bracketing:

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

Subtle though this difference may be, it distinguishes between success and failure. As far as the distributive predicate logical fragment of our language is concerned, both definitions come to the same thing. But once we add epistemic modal operators, the difference becomes important.

Given the first, global way of interpreting existential quantification $\exists x\Diamond\phi$ receives the wrong interpretation. To see that it does consider $\exists x\Diamond Px$. Using the global definition, updating with this formula would give *every* $d \in D$ as a possible value for x, as long as there is *some* d that in some world compatible with our information has the property P.

And this is clearly not what we want this formula to mean. For then success or failure of an update with $\exists x\Diamond Px$ would come to more or less the same thing as success or failure of $\Diamond\exists xPx$, the only difference being that the former has the side-effect of adding all the objects in the domain as possible values for x, whereas the latter is just a test. It would predict that $\exists x\Diamond Px$ and $\Diamond\exists xPx$ entail each other.[35]

---

ment of $\Diamond$. If it fails that test, such an assignment is eliminated. As the discussion above will have made clear, this will not yield the right empirical results.

35. Less dramatic consequences obtain if the usual definition of existential quantification is combined with a definition of the $\Diamond$-operator that is distributive with respect to the values of variables. This is the way in which Van Eijck and Cepparello proceed. This

But this is simply wrong: $\exists x \Diamond Px$ should output as possible values of x only those d such that in some w compatible with our information d has the property P in w. And this is precisely what the second, non-global definition of existential quantification does. An effect of this definition is that if $\Diamond Px$ is withing the scope of $\exists x$, the consistency test is performed one by one for each $d \in D$, where we eliminate those d as possible values for x for which the test fails.

Thus we conclude that it is the combination of interpreting $\Diamond\phi$ as global consistency testing, not only with respect to information about the world, but also with respect to discourse information, together with a non-global interpretation of $\exists x\phi$ which is distributive over the objects in the domain, that ensures that being a variable bound by and within the scope of a quantifier, differs from being a variable that is bound by, but outside the scope of a quantifier. And it is this difference that accounts for the fact that $\exists x(\phi \wedge \psi)$ and $\exists x\phi \wedge \psi$ are not equivalent precisely in those cases in which occurrences of x in ϕ or ψ are inside the scope of a $\Diamond$-operator.

4.3.3 The case of implication

We have seen in section 4.1.7 that existential quantifiers within the antecedent of an implication may bind variables in its consequent. Moreover, we noticed that if we restrict ourselves to distributive updates, $\exists x\phi \rightarrow \psi$ and $\forall x(\phi \rightarrow \psi)$ are equivalent. As was the case with extended binding, now that we have non-distributive updates at our disposal, we can give counterexamples against the generalization of the equivalence to all updates.

Fact 31
$\exists xPx \rightarrow \Diamond Qx \not\equiv \forall x(Px \rightarrow \Diamond Qx)$

This follows more or less directly from the fact that $\exists x(Px \wedge \Diamond Qx) \not\equiv \exists xPx \wedge \Diamond Qx$. Using the definition of implication, $\exists xPx \rightarrow \Diamond Qx$ can be rewritten as $\neg(\exists xPx \wedge \neg\Diamond Qx)$. Using the definition of the support test this comes down to $\neg(\exists xPx \wedge \Box\neg Qx)$. Similarly, $\forall x(Px \rightarrow \Diamond Qx)$ means the same as $\neg\exists x(Px \wedge \Box\neg Qx)$.

Consider again the situation we described in section 4.3.1. In our intial situation s_0, $\exists x(Px \wedge \Box\neg Qx)$ will only come up with the possibility

approach yield the right result for a formula such as $\exists x \Diamond Px$. However, $\exists x(Px \wedge \Diamond Qx)$ and $\exists xPx \wedge \Diamond Qx$ would still come out equivalent, which is not as it should be. Given that many other things go wrong, too.

in which the value of x is d_2 and which has w_2 as its world, which is the possibility where it is d_2 that has property P. This is so, because only for d_2 our information supports that it does not have property Q. By the definition of negation we get that $\{\langle\emptyset,\emptyset,w_1\rangle\}$ is the output of updating s_0 with of $\neg\exists x(Px \wedge \Box\neg Qx)$. And this means that we have gained the information that it is d_1 who has the property P.

Now consider $\exists xPx \wedge \Box\neg Qx$. After having updated s_0 with the first conjunct, the resulting state forbids the test $\Box\neg Qx$. This is so because there is some value for x, viz., d_1 that does have the property Qx. Hence the absurd state results after updating s_0 with $\exists xPx \wedge \Box\neg Qx$. Applying the definition of negation, this means that s_0 updated with $\neg(\exists xPx \wedge \neg\Diamond Qx)$, will result in s_0 again. No information has been gained.

This shows that there is some state in which the update with $\exists x(Px \wedge \Box\neg Qx)$ and $\exists xPx \wedge \Box\neg Qx$ give different outcomes. Hence, $\exists x(Px \wedge \Box\neg Qx) \not\equiv \exists xPx \wedge \Box\neg Qx$.

And this is as it should be. Consider the following pair of sentences:[36]

(22) Everyone who is hiding in the closet might have done it

(23) If there is someone hiding in the closet, then he might have done it

There is a subtle difference between these two sentences. Suppose we know who Alfred is. Moreover, assume that we already have the information that Alfred has not done it. Next, we are informed that (22) is the case. Subsequent investigation of the closet in question reveals that Alfred is hiding there. (Contrary to what is the case, he thought we might suspect him, and therefore he went into hiding.) We have to conclude now that the utterance of (22) was not correct.

We can also put this a follows:

Fact 32 $\neg Qa \wedge \forall x(Px \rightarrow \Diamond Qx) \wedge \Diamond Pa$ is inconsistent (in all situations in which we know who a is)

Now consider the same situation. We already have the information that Alfred has not done it. But this time, we are informed that (23) is the case. Again, we inspect the closet and we find out that Alfred is hiding there.

36. This example was brought to our attention by David Beaver. He came up with it in the discussion after a talk we gave in Amsterdam in May 1993, where we presented an earlier version of the present paper. The example was crucial in that the theory as it was then could not account for the difference between the two sentences. It was this example that led us to develop the present version of the theory.

This time we do not have to conclude that (23) was incorrect.

Fact 33 $\neg Qa \wedge (\exists x Px \to \Diamond Qx) \wedge Pa$ is consistent

Still, you may have the feeling that there is something wrong with $\neg Qa \wedge (\exists x Px \to \Diamond Qx) \wedge Pa$. What is problematic is that one cannot utter it correctly. The discourse is consistent in the technical sense, where we may assume that different parts of it are uttered with respect to different information states. But it is not consistent in the sense that we can imagine one and the same speaker believing what she says. This means that it is an incoherent discourse. It cannot be supported by any (non-absurd) information state. The reason is simply that if one believes that it is not Alfred who has done it, and at the same time believes that it is Alfred who is hiding in the closet, then of no-one who could be hiding in the closet (only Alfred is a candidate), is it compatible with our information that he has done it (Alfred has not done it).

Nevertheless, the formula $\neg Qa \wedge (\exists x Px \to \Diamond Qx) \wedge Pa$ is not incoherent. Why not? Well, one may have the information that a does not have the property Q, and that a does have the property P. But this does not exclude that there are others who also (might) have the property P, and that for at least one of them it is compatible with our information that he does have the property Q. And that is enough to believe that if someone has the property P she might have the property Q.

The point is that the discourse (23) in some sense presupposes that there is a unique culprit, and that at most one person fits in the closet. If you take those assumptions away, the discourse is no longer incoherent. (Although there still might be better ways of putting it in such a situation).

If we change the formula in such a way that we incorporate these uniqueness presuppostions with regard to the predicates P and Q (or if we make this information part of the state in which we evaluate the formula), then it does become incoherent (in that state).

4.4 Identity, coreference and modality

In this section, we continue the discussion about coreference and modality, but now we also bring identity to the fore.

Consider the following example:

(24) a. $\exists x Px \wedge \Diamond x = a \wedge x \neq a \wedge x = b$

b. Someone has done it. It might be Alfred. It is not Alfred, but it is Bill.

Obviously, this sequence of sentences, or rather its logical translation, should come out as consistent. Our semantics accounts for this fact. First notice that for (24) to be consistent in an information state, it has to meet certain preconditions. One is that there should be at least one world compatible with our information in which b has the property P, whereas a does not. Another asks that there should also be some world such that a has the property P in that world. The first requirement is that we consider it possible that Bill has done it and Alfred is innocent, the second that we do not exclude Alfred as the culprit.

Extending the situation sketched above a little, it is easy to see that the following indeed holds:

Fact 34
$\exists x Px \wedge \Diamond x = a \wedge x \neq a \wedge x = b$ is consistent

We only have to add the following to the specification of our worlds:

$$w_1(a) = w_2(a) = d_1 \;\&\; w_1(b) = w_2(b) = d_2$$

Being a is now the same property as being Q was in the original example, and being b coïncides with not being Q. This being so, the present example becomes completely similar to $\exists x Px \wedge \Diamond Qx \wedge \neg Qx \wedge Qb$. And we already saw that this formula is consistent in state s_0. Hence, it is consistent, hence $\exists x Px \wedge \Diamond x = a \wedge x \neq a \wedge x = b$ is consistent.

In the previous case, after having updated with the sequence, we know who has done it. It is the object d_2. It should be noted that updating with (24) need not always lead to a state in which we actually know who has done it, not even in case we hold on to the assumption that there is only one culprit. In the situation sketched above, we know who a and b are, they have the same denotation in each of our two worlds. But that is not necessarily so, we may very well add a couple of worlds to the information state, and cancel that knowledge.

$$w_1(P) = w_3(P) = \{d_1\} \;\&\; w_2(P) = w_4(P) = \{d_2\}$$

$$w_1(Q) = w_2(Q) = w_3(Q) = w_4(Q) = \{d_1\}$$

$$w_1(a) = w_2(a) = d_1 \;\&\; w_3(a) = w_4(a) = d_2$$

$$w_1(b) = w_2(b) = d_2 \;\&\; w_3(b) = w_4(b) = d_1$$

Now we do still know that either d_1 or d_2 has property P, and that only d_1 has property Q, and we also still know that a and b are different individuals. But we do no longer know which of d_1 and d_2 is named by a or b.

After having updated an initial situation which contains this information about the world with (24), we have gained information, because we in one sense we do know now who has done it, it is b. But in another sense we did not get any further, we still do not know who has done it, because we do not know who b is. We cannot arrest anybody yet. In this particular case, there is still hope for justice being done, by trying to find out whether b has the property Q or not. Once we get an answer to that, we will know in the stronger sense who has done it.

If however, no such distinctive properties are available, we would never be able to arrest the one who did it. We need some property, or combination of properties, that identifies d_1 or d_2, the actual objects. Or someone should be able to tell us: 'This is b', pointing at the actual person. (More about this later when it will be argued that *this* is really what we need.)

But what if we do know who is who? Consider this: you know who you are. Let us suppose you are object d_1. Hence, since you are the hearer, in every possibility in your information state, the expression 'you' refers to you, i.e., to d_1. Now consider the following sequence of sentences:

(25) a. $\exists x Px \wedge \Diamond x = you \wedge \Diamond x \neq you$
b. Someone has done it. It might be you. But it might also not be you.

We guess you also want this to be a consistent discourse. If it is your information state we are updating, then, for this to be possible, we have to imagine that you yourself are not sure about who has done it, and that moreover, you do not exclude that is was you. But this can be done. The entire affair might be something innocent, like loosing the one set of keys to the apartment you share with someone else, Or it might be something less innocent. In either case, you will be pretty keen that (25) comes out as consistent.

Well, you get away with it, but thanks only to the fact that $\Diamond\phi$ is interpreted as a global consistency test, not only with respect to information about the world, but also with respect to discourse information. If not, then after the first two sentences, your information state supports that it is you who has done it, and hence updating with the last sentence would lead to the absurd state.[37]

Let us explain. According to your information either you, d_1, or your roommate, d_2 has done it. This gives you two possible worlds. So, the first sentence of (25) does not tell you anything new. It does introduce a new peg, onto which to hang the one who has done it. In the world in which you are guilty, you hang on it, in the other world it is your roommate. Now comes the second sentence. If we test globally, the test just succeeds, and nothing happens. And the same holds for the last sentence. Everything remains possible. However, if we test the second sentence distributively, i.e., possible value by possible value, then only you pass the test, not your roommate. Since *you* refers rigidly to you, only you might be you, not your roommate. So, after updating with the second sentence in this fashion, you are the only one that hangs on the peg which was reserved for the culprit. Of course, you do not want it to be the case that merely saying that you might have done it, makes you guilty.[38] Hence you really should insist on a semantics for $\Diamond\phi$ in which it is a global test, not only with respect to information about the world, but also with respect to discourse information.

But that is not enough: even if $\Diamond\phi$ is a global test, we have a sneaky way of doing you in. We do it as follows, by saying:

(26) a. $\exists x Px \wedge \exists y(x = y \wedge \Diamond y = you)$

b. Someone has done it. It is someone who might be you.

The second peg, which is introduced by the second quantifier, in each possibility gets set to exactly the same value as the first peg, thanks to the identity. Now, the test whether it might be you is inside the scope of the second quantifier. This means that it is performed seperately for each of the possible values of the peg this quantifier has introduced, For your roommate as a value the test fails, for you it succeeds. Only you remain as a possible value for the peg introduced by the second quantifier, and hence also for the peg introduced by the first. But this implies that your information state will support that you have done it. So, you better refuse to update with this sequence.

The formulae (26a) and (27a) entail each other, but they are not equivalent. Formula (27a) is equivalent with (27b), they always lead to the

37. And this is a life sentence, once in the absurd state, you never get out anymore. Any update would lead to the absurd state again. This is actually what you are sentenced to, without any evidence whatsoever, by Van Eijck and Cepparello.

38. And neither does your roommate: if we change the order of the last two sentences, it is she that gets convicted in this all too easy way.

same update.

(27) a. $\exists x(Px \wedge \Diamond y = you)$
b. $\exists xPx \wedge \neg\neg\exists y(x = y \wedge \Diamond y = you)$

Formulae (26a) and (27a), or (27b), do not constitute one and the same update, but only because (26a) introduces two pegs, whereas (27a) and (27b) introduce only one. However, the two pegs that (26a) introduces are completely interchangeable. In every possibility they have the same value. It makes no difference whatsoever whether we continue (26a) with a formula containing a free x or a free y.

The little trick used in (26) to force a $\Diamond\phi$ test to be distributive over the values of a peg introduced by a quantifier, even when it is not in the scope of that quantifier, can be used quite generally. For example, instead of the single sentence discourse (20) given in the previous section, we can use the following two sentence discourse:

(28) a. $\exists xPx \wedge \exists y(y = x \wedge \Diamond Qy)$
b. There is someone who has the property P. It might be someone who has the property Q.

The previous example is just a special instance of this, where the property Q is the property of being you.

We end this part of the discussion with one more example. The limit case which demonstrates the need to interpret the $\Diamond$-operator globally, not only with respect to our information about the world, but also with respect to our discourse information, is where information about the world plays no role whatsoever, i.e., where the consistency test can only concern the possible values of variables. The following mathematical example is a case in point:

(29) $\exists x(x^2 = 4) \wedge \Diamond x = 2 \wedge \Diamond x = -2$

Of course, updating with the first conjunct adds no information about the world. We all know already that 4 is a square. But we do get some discourse information. A peg is introduced, and since we also know that 4 is the square of both 2 and -2, and of no other number, both 2 and -2 appear as possible values of the new peg in our information state. Information about the world plays no role here whatsoever. It is not that in half of the possible worlds the value of x is 2, whereas in the other half we have that $x = -2$. In each and every world both values are possible at the same time. So with respect to every possible world in our information state there come to correspond two possibilities, one in which the value of the new peg is 2, and one where it is -2. Actually, we might just as well assume that we have complete information about the world. In that case only one possible world is left, but there are still two possibilities according to our information: the value of x can be 2, and it can be -2.

After updating with the second conjunct, $\Diamond x = 2$, nothing should have happened to our information state. Of course, 2 is there as a possible value. And after the third conjunct, $\Diamond x = -2$, we still should be in the same state. In other words, the information state of anyone who has had some basic training in arithmetic supports (29).

Fact 35
$\models \exists x(x^2 = 4) \wedge \Diamond x = 2 \wedge \Diamond x = -2$

However, were we to perform the consistency test separately for each value of the introduced peg, then $\Diamond x = 2$ would eliminate all possibilities in which x has another value than 2. In other words, after updating with the second conjunct, our information state would support that $x = 2$. But that would mean that if we try to update with the third conjunct, $\Diamond x = -2$, the test would fail. We would end up in the absurd state. Hence, it would be predicted that (29) is inconsistent.

This is surely wrong. And that it is, shows once more that $\Diamond\phi$ should be interpreted as a global consistency test, not only with respect to information about the world, but also with respect to discourse information.

4.5 Identity and identification

In this section we explore some of the consequences of the dynamic perspective for questions of identity and identification.

For the sake of convenience we introduce the following notation convention:

Definition 38
$\exists! x\phi =_{\mathrm{df}} \exists x\phi \wedge \forall y([y/x]\phi \rightarrow x = y)$

Updating with a formula of the form $\exists! xPx$ leads to a situation in which we have the information that precisely one object has the propert P. However, as far as the identity of this object is concerned, this need not tell us anything. Consider:

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

We note the following fact:

Fact 36
$\exists !xPx \wedge \forall y \Diamond (x = y) \wedge \forall y \Diamond (x \neq y)$ is consistent

This might come as a surprise: if there is someone who has done it, how could there fail to be someone who is identical to him? But that is not what (30) expresses.[39] The kind of situation that guarantees the consistency of (30) is one in which, although we know that just one individual has the property P, we have no idea whatsoever about its identity.

Since being identical to a is a particular instance of a property that just one individual can have, similar situations can occur with respect to the identification of the referent of an individual constant:

(31) a. $\forall x \Diamond (x = a) \wedge \forall x \Diamond (x \neq a)$
b. Anyone might be Alfred. Anyone might not be Alfred.

A sequence such as (31) expresses that the value of a certain individual constant is unidentified. Analogously, (30) states that the value of a variable is unspecified.

This gives rise to the following definition:

Definition 39
Let α be a term, s an information state.

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

If a term α is identified in s, then s contains the information who α is, in at least some sense of *knowing who*.[40] If α is not identified in s, then there is at least some doubt about who α is. The notions of being identified in an information state, and of being an identifier, can be characterized as follows:

Fact 37

1. α is identified in s iff s supports $\exists x \Diamond (x = \alpha) \wedge \forall y (\Diamond (y = \alpha) \rightarrow y = x)$
2. α is an identifier iff $\models \exists x \Diamond (x = \alpha) \wedge \forall y (\Diamond (y = \alpha) \rightarrow y = x)$

Unlike variables, individual constants are meant to identify their referents, at least in ordinary cases. As long as we have not resolved the question who a is, our information is not complete. In some cases we may want to identify the value of a variable, but not for its own sake. For example, when we learn

39. Cf. $\exists !xPx \wedge \Diamond \forall y (y \neq x)$. This one surely is inconsistent.
40. Which are many. Cf. Boër and Lycan 199; Hintikka and Hintikka 1989.

that $\exists! x Px$, we may want to know who x is, because that would supply us with total information about the denotation of P. But in example (29), we would not wish to determine which of the two numbers x is, since necessarily there are two: we already know what the denotation of the predicate of being a square root of 4 is.

Within the class of individual constants, which as we said are meant to identify, some are identifiers of themselves. For example, the numerals identify their referents, the numbers, for anyone who has the slightest grasp of their meaning. But proper names generally are not identifiers.

It is of some importance to get a clear picture of what we can, and what we can not learn about the denotation of constants, since this will show that language cannot be purely descriptive, if it is to function as a means to identify objects. Let us try to make this somewhat clearer by first discussing a simple example.

Suppose we start with some state of ignorance. It suffices to consider a situation with a domain consisting of two individuals only. Let a and b be names for them. Further, let there be any number of predicates. Being ignorant, we have no idea about who a is and who b is, but let us assume that we have learned already that $a \neq b$. Furthermore, we assume that we have no idea about the denotations of the predicates.

What can we learn? A lot. For example, we 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 learnt all this, we seem to know who has the property P: a, and no-one else. About the property Q we know that it applies both to a and to b. Further, we have the information that a stands to both himself and to b in the relation R. We may imagine that we have gained such information about all the predicates.

Our knowledge is not confined to the denotations of the predicates, we also seem to know a lot about a and b. We know 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 we have learned a lot about b. Lots and lots of possibilities that our initial state of ignorance allowed, have been eliminated. Actually, with respect to some fixed set of predicates and constants, we may imagine that we have learned anything that there is to learn in this way.

But even then there are still two things we do not know. And because of that, there are lots of other things we do not know either. Our 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 we know that one individual has that property, but we have no idea who this is. With respect to Q things are different: since we know that both a and b have Q, we are certain as to who is Q: everyone. As for R, there is again uncertainty. We know which pairs form its extension, but since these are not all pairs, and since we do not know who a and b are, there is a sense in which we do not know between which individuals the relation holds.

So, although we have learned all there is to learn in this way, we have not, and will not come to know who is a and who is b. We can formulate this as follows.

Definition 40
Let $\langle r, g, w\rangle, \langle r, g', w'\rangle \in I$
$\langle r, g, w\rangle \simeq \langle r, g', w'\rangle$ iff there exists a bijection f from D onto D such that:
1. For every peg m in the domain of g: $g'(m) = f(g(m))$
2. For every individual constant a: $w'(a) = f(w(a))$.
3. 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 38
For all $s \in S$ such that $s = \mathbf{0}[\phi_1]\ldots[\phi_n]$: if $i \in s$ then $i' \in s$, for every $i' \simeq i$

Roughly speaking, what this fact tells us is that there is no way that we can identify individuals by purely descriptive means. If we start out not knowing who is who, no matter with how many descriptive sentences we update our information state, the result will never allow us to really identify an individual. It will always be possible to permute individuals without affecting the information that we have, and hence the sentences that it supports (forbids, allows).

At first sight this may seem surprising, but on second thought (sic!) it is really a perfectly normal result. No matter how important linguistic communication is for obtaining information about the world, it can not be all that there is to it: somehow we need to anchor it, by essentially non-linguistic means, onto real objects and properties.

In order words, in order to use linguistically conveyed information to identify individuals, we also at some point need to go out in the world, and have a look. One way this works is through observation of the instantiations of properties. At least some of our predicates will be linked with observational criteria. For example, we may know what objects that have the property P look like. This we can make use of. We have a look, we inspect object d_1, and satisfy ourselves that it meets the observational criteria, and hence has the property in question. We know: 'This [*pointing at* d_1] is a'. In our example this is actually enough to give us total information: there being only two objects, identifying a means identifying b, and hence we now really know everything. We know know who a is, who b is, who has the property P, and who stand in the relation R.

To satisfy our need for identification, we introduce demonstratives into our language.

Definition 41
We add *this* to the logical inventory of the language.

- If $d \in D$, then $this_d$ is a term.
- For all $i \in I$: $i(this_d) = d$

This is, of course, a rather crude way of dealing with demonstratives, but for our present purposes it suffices. The general idea is that which object is pointed at when *this* is uttered, is part of the discourse information.[41] We will leave out the explicit indication of what object is pointed at whenever this is convenient. Also, next to *this*, we use *that*. So, *this* and *that* are two possibly different demonstrative identifiers.

Fact 39

1. $this_d$ is an identifier
2. $\models \Diamond(this = that) \rightarrow (this = that)$
3. $\models (this = that) \rightarrow \Box(this = that)$

Suppose there are just two different things around that we can point at. And suppose we know that one of them is called a, and the other b, but that we do not know which is called which. In such a situation we have that:

$$(this \neq that) \wedge \neg\exists x((x \neq this) \wedge (x \neq that))$$

$$(a \neq b) \wedge \neg\exists x((x \neq a) \wedge (x \neq b))$$

Our information state also supports the following sequence of tests:

$$\Diamond(this = a) \wedge \Diamond(that = a) \wedge \Diamond(this = b) \wedge \Diamond(that = b)$$

41. It would be more elegant to let demonstratives (and invidual constants) introduce a peg, and to introduce a function in the referent system that gives us the object pointed at. Things can be set-up in that way, but since the added complexity does not pay off here, we do not bother. Should we need pegs for constants and demonstratives, we can always obtain them in the present set-up by introducing these expressions by means of existential quantification: $\exists x(x = a)$, $\exists x(this_d = x)$.

Again, these results are what one would like them to be in an epistemic setting.

We conclude this section by pointing out two specific logical properties of the present treatment. Let the situation be as we just described it. In such a situation someone might be inclined to reason as follows. Since a can be *this* and b can be *this*, and since there is nothing else around, why not conclude:

$$\forall x \Diamond (x = this)$$

However, no such conclusion is warranted, because a and b are not identifiers. Quantification is over *objects*, and a and b do not identify objects. But *this* and *that* do, so the following reasoning *is* correct Since *this* can be a, and *that* can be a, and *this* and *that* are all there is, a can be anything:

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

At this juncture we would again take a wrong turn if we would conclude that since everything can be a, and since b is something, we therefore have that:

$$\Diamond (b = a)$$

This contradicts the assumption that $a \neq b$, and hence this line of reasoning would result in absurdity. Of course, b is something, but in our epistemic setting that is not enough to allow the instantiation from $\forall x \Diamond (x = a)$ to $\Diamond (b = a)$. We need a a stronger premise here, viz., that there is something that *must* be b. And that our information does not guarantee.

We conclude that universal instantiation is not always valid. For similar reasons existential generalization does not hold unconditionally either:

$$\forall y \Diamond y \neq a \not\models \exists x \forall y \Diamond y \neq x$$

However, for identifiers both logical principles do hold:

Fact 40
1. If α is an identifier, then $\forall x \phi \models [\alpha/x]\phi$
2. If α is an identifier, then $[\alpha/x]\phi \models \exists x \phi$

These few remarks may have given the reader some idea of the many interesting issues that are involved in a dynamic treatment of identity. Other questions have to do, for example, with the notion of specificity, the relation between epistemic and metaphysical identity, and so on. For the larger part these are traditional questions, but we hope that the few observations reported in this section, have shown that studying these in a dynamic framework puts them in a new perspective.

5 Discourse coherence

An observation we have made repeatedly in previous sections, is that some sequences of sentences, e.g. $\Diamond p \wedge \neg p$, can be consistent, and at the same time incoherent. We explained the latter as meaning that such sequences of sentences are not acceptable when interpreted as a single utterance by a single speaker. However, as a multi-speaker discourse, or if we imagine that some time elapses between the utterances of the sentences in the sequence, allowing for new information coming in from outside, such sequences can constitute an acceptable coherent discourse. This introduces a notion of discourse coherence next to the notion of sentence coherence, that is the topic of this final section.

Let us return for a moment to example (29) discussed in the previous section, which we repeat below:

(29) $\exists x(x^2 = 4) \wedge \Diamond x = 2 \wedge \Diamond x = -2$

In some respect this mathematical example differs from the others that we have discussed. Consider the following continuation:

(32) $\exists x(x^2 = 4) \wedge \Diamond(x = 2) \wedge \Diamond(x = -2) \wedge (x = 2)$

In any intuitive sense of the word, this sequence of formulae is unacceptable. If you are not convinced right away,[42] you will be in a minute.

In what way does (32) differ from the sequence (33), which, as we have seen, is acceptable?

(33) $\exists x Px \wedge \Diamond(x = a) \wedge \Diamond(x = b) \wedge (x = a)$

There are several differences. One is that it need not be the case that a and b are identified in our information state, i.e., we need not know to which

42. If you are not, then for you the language of modal predicate logic is not a 'natural' language, and, probably, what you do is that you immediately start translating it into a more familiar vernacular. With this particular example, it is important to translate it quite 'literally', i.e., as something like:

(i) There is at least one number the square of which is four. It might be two, it might be minus two. It is two.

You see! But it is even better to stick to the formula as such, rather than consider its translation. After all, despite its 'unnaturalness' a logical language is also a language you can have intuitions about, once you have learned the basics of its semantics.

individuals a and b refer. For 2 and -2 this is different, they rigidly refer to the numbers 2 and -2, respectively. That is not the difference which we want to draw attention to, however. Let us therefore abstract from this difference, and compare (32) with:

(34) $\exists x Px \wedge \Diamond(x = 2) \wedge \Diamond(x = -2) \wedge (x = 2)$

The only difference now is that wheras in (34) we have just some property P, in (32) the property of being a square root of 4 is at stake. To make things more concrete, we think of P as some contingent property, e.g., the property of being the number the teacher wrote on the back of the blackboard. In that case (34) corresponds to the following sequence of sentences:

(35) The teacher has written a number on the back of the blackboard. It might be 2, it might be -2. It is 2.

This is also pretty odd, but in a different way than (32) is. Above we have explained this kind of deviance: (35) is unacceptable as a one-speaker discourse, because it is incoherent. But this discourse becomes completely in order if we assume it is a multi-speaker discourse. For example we can imagine a situation in which the first two sentences are uttered by pupil A, sitting in the class-room, and the third one comes from a fresh boy, B, who has managed to sneak up to the front of the class-room and has seen the number which the teacher actually wrote down.

(36) A: The teacher has written a number on the back of the blackboard. It might be 2, it might be -2. B: It is 2.

In this situation (35) is alright, and that is exactly why it can be acceptable, even though it is incoherent as a one speaker discourse. For as such, that does not exclude that it is acceptable as a multi-speaker discourse.

Let us now return to our example (32). This sequence is equally consistent, and equally incoherent as a one-speaker discourse. But unlike in the previous case, constructing it as a multi-speaker discourse does not improve it:

(37) A: $\exists x(x^2 = 4) \wedge \Diamond(x = 2) \wedge \Diamond(x = -2)$ B: $x = 2$

How should this be explained? No matter how sneaky B is, he can never be in an information state which supports $x = 2$, *if* he has updated his information with what A has said. And he must have. This is signalled by his use of the anaphoric pronoun *it* to refer back to the peg introduced by the quantifier in the first sentence that A has uttered.[43] If B has had some

43. Such anaphoric links 'across speakers' provide an interesting challenge for existing theories of anaphora.

basic training in arithmetic, then he knows that in every possible world, 4 has two square roots: 2 and -2. So both are possible values of x with respect to every world that is still possible according to his information. Hence, he has no support for his utterance that $x = 2$.

Notice that (32) *is* acceptable if we assume that B has updated only with the first sentence uttered by a, and has rejected the second, for example, because he, erroneously, believes that 4 has only one square root, viz., 2. In such a situation B's utterance , although supported by his information state, would still be incorrect, in this sense, that not signalling the discrepancy between what one believes and what other speech participants say, is not co-operative.

This, then, is the difference between (32) and (34): the property of being a square root of 4 rigidly denotes the set $\{2, -2\}$, the property of being written at the back of the blackboard does not.

A formal account of these intuitive explanations requires a notion of discourse coherence. As such coherence is a very important and central, but also (therefore?) a very complicated one. The following definitions should be interpreted as a provisional first step towards an account of one aspect of this complicated matter.

First of all, it may be convenient at this point to make a formal distinction between a discourse and a conjunction of sentences. And we assume that different sentences in a discourse may be uttered by different speakers. Above we saw that in a two-speaker discourse one of the requirements is that a speech participant must either be in a state that allows the utterance of the other participant and then updates with it, or, if his information forbids the utterance of the other, then he must signal this, leaving his information state untouched.

In the following definition we introduce discourses, and define the effect of consequtively updating an information state with the sentences in a discourse. Discourse update is defined in such way that in updating with the sentences in a discourse, we 'skip' the ones that would lead to the absurd state.

Definition 42

Let $s \in S$.

1. If $\phi_1 \ldots \phi_n$ are formulas, then $\phi_1; \ldots\ ; \phi_n$ is a discourse
2. $s[\![\phi]\!] = s[\phi]$ if s allows ϕ
   $s[\![\phi]\!] = s$ if s forbids ϕ
3. $s[\![\phi; \psi]\!] = s[\![\phi]\!][\![\psi]\!]$

Each participant updates with all the sentences in the discourse, also the ones he utters himself. We will also require that his information state supports his own utterances. This means that updating with his own utterances cannnot give new 'real' information about the world, but such updates can provide new discourse information, if a quantifier is introduced.

Although in the interpretation of a discourse we just 'skip' utterances that are inconsistent with our information, we have to indicate to the other participant(s) that this special situation has occurred. Signalling a conflict between what you believe and what someone else says, can be done in very many ways. For our present purposes, a convenient device is that of uttering something that simply contradicts what is being said. The relevant relation is defined as follows:

Definition 43
ϕ *contradicts* ψ iff $\forall s \in S$ such that $s \neq \emptyset$: s supports $\psi \Leftrightarrow s$ forbids ϕ

As is to be expected, $\neg\phi$ contradicts ϕ. But also, e.g., $\neg\phi$ contradicts $\Diamond\phi$.

We first define a notion of a coherent exchange of two sentences between two speech participants. Then we use that to define what a coherent dialogue is. We assume: speakers should support what they say. Hearers either allow what is said and update, or respond with an explicit rejection:

Definition 44
Let $s, s' \in S$.
$\phi;\psi$ is a *coherent exchange* for $\langle s, s' \rangle$ iff

1. s supports ϕ
2. s' allows ϕ and $s'[\phi]$ supports ψ; or
   s' forbids ϕ and ψ contradicts ϕ

Let A utter ϕ, and B reply with ψ. For this exchange to be coherent it is first of all required that the information state of A support ϕ, i.e., that A believe that ϕ is the case. As for B, if his state is such that he is able to take in A's utterance, he may continue with some other utterance ψ, provided his information supports it. However, if B's information state forbids ϕ, he is to signal this, by uttering some sentence ψ which contradicts ϕ.

If we generalize to a dialogue of arbitrary length, we must assure that the two speech particants take turns adequately:

Definition 45

Let $s_1, s_2 \in S$.
$\phi_1; \ldots; \phi_n$ is a *coherent dialogue* for $\langle s_1, s_2\rangle$ iff

1. For all k: $1 \leq k < n$:
   ϕ_k, ϕ_{k+1} is a coherent exchange for $\langle s[\![\phi_1; \ldots; \phi_{k-1}]\!], s'[\![\phi_1; \ldots; \phi_{k-1}]\!]\rangle$, where $s = s_1, s' = s_2$ if k is odd, and $s = s_2, s' = s_1$ if k is even
2. For $k = n$: $s[\![\phi_1; \ldots; \phi_{n-1}]\!]$ allows ϕ_n, where $s = s_1$ if n is even, and $s = s_2$ if n is odd.

Finally, we can now say when a dialogue is coherent:

Definition 46
$\phi_1; \ldots; \phi_n$ is a *coherent dialogue* iff $\exists s, s' \in S$ such that $\phi_1; \ldots; \phi_n$ is a *coherent dialogue* for $\langle s, s'\rangle$

A dialogue is coherent if the first sentence is supported by the initial information state (of the first speaker), and if each of the following utterances is supported by an information state (of the speaker whose turn it is) which is the result of either updating with was has gone before, or rejecting it and signalling that.

Note that it is not forbidden that any of the $k-1$ information states coincide: not only the one-speaker – one-hearer case is covered, but also that of the speaker that nobody listens to, except she herself. What *is* required is that the speakers listen to each other, which is hard enough. The ultimate hearer could also be Nature, i.e., the minimal information state. But unless nothing of substance is said, the speaker or speakers cannot be Nature. Nature does not speak for itself.

According to this definition of discourse coherence we have that:

Fact 41

1. $\exists x(x^2 = 4) \wedge \Diamond(x = 2) \wedge \Diamond(x = -2) \wedge (x = 2)$ is an incoherent discourse
2. $\exists x Px \wedge \Diamond(x = a) \wedge \Diamond(x = b) \wedge (x = a)$ is a coherent discourse
3. $\exists x Px \wedge \Diamond(x = 2) \wedge \Diamond(x = -2) \wedge (x = 2)$ is a coherent discourse
4. $\neg Pa \wedge \exists x \Diamond Px \wedge (x = a)$ is an incoherent discourse

It seems that these predictions are borne out by our intuitions, if the latter are sufficently sensitive to the difference between the speaker and the hearer perspective, and to the distinction between a one-speaker and a multi-speaker discourse.

By way of illustration we discuss the last example.[44] Assume that the referent of a has not been identified yet. If it were, then the discourse would simply be inconsistent. Assuming that it is not, the discourse *is* consistent. Consider the following sequence which embodies this logical structure, and which satisfies this assumption:

(38) Maria, whoever she is, has not done it. There is someone who might have done it. It is Maria.

Intuitively, (38) is an unacceptable discourse. It is consistent, as long as Maria has not been identified, and it is incoherent, if it is taken as a one-speaker discourse. But that still does not explain its unacceptability completely. We have seen several examples of discourses that were unacceptable as one-speaker discourses, but which could be interpreted as acceptable, as soon as we assumed that more than one-speaker was involved. But in case of (38), the unacceptability seems much stronger than that. It does not matter who says what. As long as they all take part in the same conversation, from beginning to end, the result is wrong.

This property of (38) is adequately explained by the notion of coherence given above. If you take part in this dialogue, then whatever your initial state of information, if you can get along with the conversation up to the third sentence, i.e, if you did not in the meantime reject what was said before as being inconsistent with your information, then after having updated with the first two sentences, there must be some possibility i in your information state, such that $i(x) \in i(P)$, and, more importantly, $i(x) \neq i(m)$. But that possibility would be eliminated by the last sentence $x = a$. Hence, at that point you could not correctly utter $x = a$. It is not supported by your information state. (Of course, things are different if you would have missed the first sentence.) If you utter it none the same, then the other participant will refuse to update with your utterance, even if her information allows an update with it. She knows the rules of the game, hence she knows your information cannot possibly warrant you to say $x = a$, unless that is you forgot to protest against one of her previous utterances, or missed part of the conversation.

Let us end this final section with some edifying remarks about communication. The notion of discourse coherence requires nothing whatsoever about how the initial information states of the participants in a discourse relate to each other. In particular, it is not required that any two of them are

44. It was brought to our attention by Maria Alloni. The notion of discourse coherence defined in the text was the result of trying to account for this type of discourse.

compatible with each other. Actually, they can all be completely disjunct. Moreover, none of them needs to include the actual world. This is very realistic. It is more than likely that for any two participants in a discourse, there is at least one fact about which they disagree, about which one believes that it is the case, and the other that it is not. If so, there is no world that both of them consider to be possible. At the same time, it is also more than likely that each of the participants believes at least some fact to hold that does not actually hold, whence the actual world for none of them will be among their possible worlds. This is nothing to get excited about, it is just the human condition.

Still, we can talk, and we can all learn from it. As long as we avoid to talk about the facts about which we happen to disagree, our conversations can be informative for all particpants. We all may get better informed in the sense that some worlds we still considered possible, may be eliminated. For any two participants this may mean that after a litlle chat, they at least agree more on what possibilities are excluded. The dark side of this is that we may happily agree on excluding the wrong worlds, i.e., that we exclude worlds that are more like the actual one than the ones that we stick to. There is no guarantee that by agreeing more, or disagreeing less, we get any nearer to the truth.

Acknowledgements

We owe a special thanks to Paul Dekker. The present paper builds heavily on the last chapter of his thesis. The use of referent systems is inspired by the work of Kees Vermeulen. The work of Jan van Eijck and Giovanna Cepparello on the subject was another source of inspiration to us. David Beaver, Maria Alloni and Jelle Gerbrandy provided us with new data that helped us a lot.

Earlier versions of the paper were presented on various occasions in the past year. The first of these was a lengthy presentation at the Workshop on Tense and Modality (Columbus, Ohio, July 1993). Thanks to Craige Roberts for inviting us, and for taking good care of us. For their helpful comments we thank the participants of that workshop, and of other events where we talked about this material.

An extended abstract of this paper will appear in D. Westerståhl & J. Seligman, *Language, Logic and Computation: The 1994 Moraga Proceedings*, Stanford: CSLI, under the title 'This Might Be It'.

Bibliography

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

Beaver, D. I.: 1993a, Kinematic Montague grammar, in H. Kamp (ed.), *Presupposition*, ILLC, University of Amsterdam, Amsterdam, Dyana-2 deliverable R2.2.A, Part II

Beaver, D. I.: 1993b, The kinematics of presupposition, in H. Kamp (ed.), *Presupposition*, ILLC, University of Amsterdam, Amsterdam, Dyana-2 deliverable R2.2.A, Part II

Beaver, D. I.: 1993c, Two birds and one stone, in H. Kamp (ed.), *Presupposition*, ILLC, University of Amsterdam, Amsterdam, Dyana-2 deliverable R2.2.A, Part II

Benthem, J. v. and Cepparello, G.: 1994, *Tarskian Variations. Dynamic parameters in classical semantics*, Technical Report CS-R9419, CWI

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.

Bunt, H.: 1990, Dynamic interpretation in text and dialogue, in L. Kálmán and L. Pólos (eds.), *Papers from The Second Symposium on Logic and Language*, pp 67–104, Akadémiai Kiadó, Budapest

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

de Vrijer, R. C.: 1990, *Natural Deduction for DPL*, unpublished manuscript, Department of Philosophy, University of Amsterdam

Dekker, P.: 1988, Predicate logic with anaphora (seven inch version), in L. Santelmann and M. Harvey (eds.), *Proceedings of the Fourth SALT Conference*, DMLL Publications, Cornell University

Dekker, P.: 1990, Dynamic interpretation, flexibility and monotonocity, in M. Stokhof and L. Torenvliet (eds.), *Proceedings of the Seventh Amsterdam Colloquium*, ITLI, Amsterdam

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

Dekker, P.: 1994, Predicate logic with anaphora, in R. H. Cooper and J. A. Groenendijk (eds.), *Integrating Semantic Theories II*, ILLC, Amsterdam, Dyana-2 deliverable R2.1.B

Fernando, T.: 1994, Generalized qantifiers as second-order programs —

'dynamically' speaking, naturally, in M. Stokhof and P. Dekker (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 249–68, ILLC, University of Amsterdam, Amsterdam

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

Gärdenfors, P.: 1988, *Knowledge in Flux: Modelling the Dynamics of Epistemic States*, MIT Press, Cambridge Mass.

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.: 1994, Modality and coreference, in R. H. Cooper and J. A. Groenendijk (eds.), *Integrating Semantic Theories II*, ILLC, Amsterdam, Dyana-2 deliverable R2.1.B

Groenendijk, J., Stokhof, M., and Veltman, F.: to appear, Pronouns and scope, in S. Lappin (ed.), *Handbook of Contemporary Semantic Theory*, Blackwell, Oxford

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

Groeneveld, W. and Veltman, F.: 1994, Inference systems for update semantics, in R. H. Cooper and J. A. Groenendijk (eds.), *Integrating Semantic Theories II*, ILLC, Amsterdam, Dyana-2 deliverable R2.1.B

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

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

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

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

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

Israel, D.: 1994, The very idea of dynamic semantics, in M. Stokhof and P. Dekker (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 1–16, ILLC, University of Amsterdam, Amsterdam

Janssen, T. M.: 1990, Models for discourse markers, in M. Stokhof and L. Torenvliet (eds.), *Proceedings of the Seventh Amsterdam Colloquium*, ITLI, University of Amsterdam, Amsterdam

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

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

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

Kanazawa, M.: 1993, *Dynamic Generalized Quantifiers and Monotonicity*, Prepublication LP-93-02, ILLC, University of Amsterdam, Amsterdam

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

Krahmer, E.: 1994, Partiality and dynamics, theory and application, in P. Dekker and M. Stokhof (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 391–410, ILLC, University of Amsterdam, Amsterdam

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. e. a. Egli (ed.), *Interface Aspects of Syntax, Semantics and the Lexicon*, Benjamins, Amsterdam/Philadelphia

Pagin, P. and Westerståhl, D.: 1994, Flexible variable binding and Montague grammar, in P. Dekker and M. Stokhof (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 519–27, ILLC, Amsterdam

Pagin, P. and Westerståhl, D.: to appear, Predicate logic with flexibly binding operators and natural language semantics, *Journal of Logic, Language and Information*

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

Roberts, C.: 1989, Modal subordination and pronominal anaphora in dis-

course, *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
Stalnaker, R.: 1979, Assertion, in P. Cole (ed.), *Syntax and Semantics 9 – Pragmatics*, Academic Press, New York
Stalnaker, R.: 1984, *Inquiry*, MIT Press, Cambridge, Mass.
van Benthem, J.: 1989, Semantic parallels in natural language and computation, in H. Ebbinghaus et al. (eds.), *Logic Colloquium '87*, Elsevier, Amsterdam
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 Benthem, J. and van Eijck, J.: 1982, The dynamics of interpretation, *Journal of Semantics* 1(1), 1–20
van den Berg, M. H.: 1990, A dynamic predicate logic for plurals, in M. Stokhof and L. Torenvliet (eds.), *Proceedings of the Seventh Amsterdam Colloquium*, ITLI, University of Amsterdam, Amsterdam
van den Berg, M. H.: 1993, Dynamic generalized quantifiers, in J. van der Does and J. van Eijck (eds.), *Generalized Quantifier Theory and Applications*, The University of Chicago Press, Chicago, Ill.
van den Berg, M. H.: 1994, A direct definition of generalized dynamic qantifiers, in M. Stokhof and P. Dekker (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 121–40, ILLC, University of Amsterdam, Amsterdam
van der Does, J.: 1993, The dynamics of sophisticated lazines, in J. Groenendijk (ed.), *Plurals and Anaphora*, ILLC, University of Amsterdam, Amsterdam, Dyana-2 deliverable R2.2.A, Part I
van der Does, J.: 1994a, Formalizing e-type anaphora, in M. Stokhof and P. Dekker (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 229–48, ILLC, University of Amsterdam, Amsterdam
van der Does, J.: 1994b, *'Might be'. Some Proof Theory for a Substructural Dynamic Logic*, unpublished manuscript
van Eijck, J.: 1993, The dynamics of description, *Journal of Semantics* 10(2)
van Eijck, J.: 1994a, The dynamics of theory extension, in M. Stokhof and P. Dekker (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 249–68, ILLC, University of Amsterdam, Amsterdam
van Eijck, J.: 1994b, *Presupposition Failure: A Comedy of Errors*, unpub-

lished manuscript, CWI, Amsterdam. To appear in *Methods of Logic in Computer Science*

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.: 1992a, Dynamic interpretation and Hoare deduction, *Journal of Logic, Language and Information* 1(1), 1–44

van Eijck, J. and de Vries, F.-J.: 1992b, A sound and complete calculus for dynamic predicate logic, in P. Dekker and M. Stokhof (eds.), *Proceedings of the Eighth Amsterdam COlloquium*, ILLC, University of Amsterdam, Amsterdam

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.: 1990, Defaults in update semantics, in H. Kamp (ed.), *Conditionals, Defaults, and Belief Revision*, CCS, Edinburgh, Dyana deliverable R2.5.A

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

Verkuyl, H. J. and Vermeulen, C. F.: 1993, *Shifting Perspectives in Discourse*, unpublished manuscript, OTS, Utrecht

Vermeulen, C. F.: 1993a, *Proofs as Texts. Dynamic Proof Theory for Intuitionistic Propositional Logic*, Technical Report OTS-WP-CL-93-007, OTS, Utrecht

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

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

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

Vermeulen, C. F.: to appearb, 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
Zeevat, H.: 1994, An update semantics for questions, in H. Kamp (ed.), *Ellipsis, Tense and Questions*, Dyana-2 deliverable R2.2.B