This Might Be It

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

September 1994

1 Introduction

Discussions often end before the issues that started them have been resolved. For example, in the late sixties and early seventies, a hot topic in philosophical logic was the development of an adequate semantics for the language of modal predicate logic. However, the result of this discussion was not one single system that met with general agreement, but a collection of alternative systems, each defended most ably by its proponents.

Although it would seem that the topic has lost much of its controversial status, this paper adds one more system to the existing stock. It offers a semantics for the language of modal predicate logic, which is new, not in the sense that it proposes a new ontology as an alternative to the possible world paradigm, but new because it characterizes the meaning of a sentence in terms of its information change potential rather than its truth conditions. What we hope to show is that this dynamic twist sheds new light on old

*. We owe a special thanks to Paul Dekker. The present paper builds heavily on the last chapter of his thesis. His comments on various stages of the work reported here have prevented us from making many mistakes. For the remaining ones we take the blame. Maria Alloni and Jelle Gerbrandy also provided useful feed-back. Earlier versions of the paper were presented on various occasions. The first of these was the Workshop on Tense and Modality (Columbus, Ohio, July 1993). For their helpful comments, we thank the participants of that workshop, and of other events where we talked about this material. The paper will appear in D. Westerståhl & J. Seligman, *Language, Logic and Computation: The 1994 Moraga Proceedings*, Stanford: CSLI. It is an extended abstract of a longer paper, 'Update semantics for modal predicate logic', by the same authors.

issues concerning modality, (co)reference, identity and identification.[1]

The idea that meaning is information change is implemented by interpreting sentences as updates, as functions from information states to information states. From this epistemic perspective, the notion of truth, which relates language to the world, loses its key role. Central notions are consistency and support, which relate language, not to the world, but to the information language users have about it. Consequently, the entailment relation has to face change, too.

Since interpretation is viewed as a process of updating information, the structure and contents of information states has to be explicated. This is done section 2. In section 3, an update semantics for the language of modal predicate logic is stated, which is illustrated by a discussion of some representative examples in section 4. In section 5, special attention is paid to problems of identity and identification, some of which, it is argued, necessitate the introduction of demonstratives. In section 6 we look ahead.

2 Information

An information state is looked upon as a set of possibilities, viz., those alternatives which are still open according to someone in that state. What the possibilities are depends on what the information is about. First, there is information about the world. Quite often, such information is gathered by verbal means. The interpretation of discourse raises its own questions. For example, there is the issue of resolving anaphoric relations. We have to keep track of what we talk about. This kind of discourse information is more like a book-keeping device than like real information. Yet, it is essential for the interpretation of discourse, and since the latter is an important source of information about the world, discourse information, indirectly, also provides such information.

Information about the world is represented as a set of possible worlds, those worlds that given the information available still might be the real one. Worlds are identified with first order models. In the present paper, it is assumed that the language users know which objects constitute the domain of discourse (although they may not know their names). In view of this, all

1. The present system combines dynamic predicate logic with update semantics for modal propositional logic (see Groenendijk and Stokhof 1991; Veltman 1990). On the problem of combining the two see, e.g., Groenendijk and Stokhof 1990, Dekker 1993, chapter 5, Vermeulen 1994, chapter 2.

possible worlds share one domain. Hence, a possible world can be identified with an interpretation function of a first order model. Extending information about the world amounts to eliminating worlds from the ones which were still considered possible.

The language interpreted is a logical language with quantifiers and variables. The use of a quantifier introduces a new item of conversation, a new *peg*. Variables are the anaphoric expressions of the language. To enable the resolution of anaphoric relations, discourse information keeps track, not only of the number of pegs, but also of the association between variables and pegs. Extending discourse information is adding variables and pegs.

Discourse information is linked to information about the world, *via* possible assignments of objects to the pegs (and hence, indirectly, to the variables associated with these pegs). In general, not every assignment of an object to a peg is possible — both the discourse and the available information may provide restrictions —, but usually, more than one is. Getting better informed on this score is eliminating possible assignments. Suppose a certain assignment is the only one left with respect to some world which is still considered possible. In that case elimination of the assignment brings along the elimination of the world. This is how discourse information may provide information about the world.

2.1 Information States

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

Definitie 1 A *referent system* is a function r, which has as its domain a finite set of variables v, and as its range a number of pegs.
If the number of pegs in a referent system is n, then the numbers $m < n$ are its pegs.

The use of a quantifier $\exists x$ introduces the next peg, and associates the variable x with that peg:

Definitie 2 Let r be a referent system with domain v and range n.

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

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

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

Associating a variable with a new peg is the proto-typical way in which a referent system is extended:

Definitie 3 Let r and r' be two referent systems with domain v and v', and range n and n', respectively.
r' is an *extension* of r, $r \leq r'$, iff $v \subseteq v'$; $n \leq n'$; if $x \in v$ then $r(x) = r'(x)$ or $n \leq r'(x)$; if $x \notin v$ and $x \in v'$ then $n \leq r'(x)$.

A referent system $r[x/n]$ is always a *real* extension of r.

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

Definitie 4 Let D, the *domain of discourse*, and W, the set of *possible worlds*, be two disjoint non-empty sets.
The *possibilities* based on D and W is the set I of triples $\langle r, g, w\rangle$, where r is a referent system; g is a function from the range of r into D; $w \in W$.

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

Information states are (real) subsets of the set of possibilities:

Definitie 5 Let I be the set of possibilities based on D and W.
The set of *information states* based on I is the set S such that $s \in S$ iff $s \subseteq I$, and $\forall i, i' \in s$: i and i' have the same referent system.

Variables and pegs are introduced globally with respect to information states. That is why an information state has a unique referent system.

2.2 Information Growth

One way in which information can grow is by adding variables and pegs and assigning some object to them:

Definitie 6 Let $i = \langle r, g, w\rangle \in I$; n the range of r; $d \in D$, $s \in S$.
1. $i[x/d] = \langle r[x/n], g[n/d], w\rangle$
2. $s[x/d] = \{i[x/d] \mid i \in s\}$

Information can grow, not just by adding discourse information, but also by eliminating possible assignments of objects to pegs, and by eliminating possible worlds.

Definitie 7 Let $i, i' \in I, i = \langle r, g, w\rangle$ and $i' = \langle r', g', w'\rangle$, and $s, s' \in S$.
1. i' is an *extension* of i, $i \leq i'$ iff $r \leq r'$, $g \subseteq g'$, and $w = w'$
2. s' is an *extension of* s, $s \leq s'$ iff $\forall i' \in s' \colon \exists i \in s \colon i \leq i'$

The extension relation is a partial order. There is a unique minimal information state, the *state of ignorance* in which all worlds are still possible and no discourse information is available yet: $\mathbf{0} = \{\langle \emptyset, \emptyset, w\rangle \mid w \in W\}$. The maximal element in the ordering $\mathbf{1} = \emptyset$, the *absurd state* in which no possibility is left. Less maximal, but more fortunate, are states of *total information*, consisting of just one possibility.

Some auxiliary notions:

Definitie 8 Let $s, s' \in S, s \leq s', i \in s, i' \in s'$.
1. If $i \leq i'$, we say that *i' is a descendant of i in s'*
2. If i has one or more descendants in s', we say that *i subsists in s'*
3. If all $i \in s$ subsist in s', we say that *s subsists in s'*

3 Updating Information States

Now that information states are defined, they can be put to use: in updating. The formulae of the familiar language of modal predicate logic are interpreted as (partial) functions from information states to information states.

We use postfix notation: $s[\phi]$ is the result of updating s with ϕ, $s[\phi][\psi]$ is the result of first updating s with ϕ, and next updating $s[\phi]$ with ψ. Whether s can be updated with ϕ may depend on the fulfillment of certain constraints. If a state s does not meet them, then $s[\phi]$ does not exist, and the interpretation process comes to a halt.

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

Definitie 9 Let α be a basic expression, $i = \langle r, g, w \rangle \in I$, with v the domain of r, and I based upon W and D.

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 $i(\alpha)$ is not defined
3. If α is an n-place predicate, then $i(\alpha) = w(\alpha) \subseteq D^n$

The absence of a variable in the referent system of a state will be the only source of partiality of updates.

The following definition specifies the update semantics for the language of modal predicate logic:

Definitie 10

1. $s[Rt_1 \dots t_n] = \{i \in s \mid \langle i(t_1), \dots, i(t_n) \rangle \in i(R)\}$
2. $s[t_1 = t_2] = \{i \in s \mid i(t_1) = i(t_2)\}$
3. $s[\neg\phi] = \{i \in s \mid i \text{ does not subsist in } s[\phi]\}$
4. $s[\phi \wedge \psi] = s[\phi][\psi]$
5. $s[\exists x\phi] = \cup_{d \in D}(s[x/d][\phi])$
6. $s[\Diamond\phi] = \{i \in s \mid s[\phi] \neq \emptyset\}$

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

Notice that atomic updates can be partial. If one of the argument terms is a variable that is not present in the referent system of the information state, then its denotation is not defined, and hence the update does not exist. This carries over to all the other update clauses. If somewhere in the interpretation process we meet a variable that at that point has not been introduced, then the whole process comes to a halt.

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

Updating a state with a conjunction is a sequential operation: the state is updated with the first conjunct, and next the result of that is updated with the second conjunct. The interpretation of a conjunction is the composition of the update functions associated with the conjuncts.

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

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

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

Feit 1 For every formula ϕ and information state s: $s \leq s[\phi]$

The observation tells that we are justified in calling the semantics an *update semantics*. The interpretation process always leads to an information state that is an extension of the initial state.

Other logical constants can be added by definition in the usual way. Calculating the definitions out, we get:

Feit 2

1. $s[\phi \rightarrow \psi] = \{i \in s \mid$ if i subsists in $s[\phi]$, then all descendants of i in $s[\phi]$ subsist in $s[\phi][\psi]\}$
2. $s[\phi \vee \psi] = \{i \in s \mid i$ subsists in $s[\phi]$ or i subsists in $s[\neg\phi][\psi]\}$
3. $s[\forall x\phi] = \{i \in s \mid$ for all $d \in D$: i subsists in $s[x/d][\phi]\}$
4. $s[\Box\phi] = \{i \in s \mid s$ subsist in $s[\phi]\}$

It should be remarked that it is not possible to make a different choice between basic and defined operations that leads to the same overall results.

3.1 Consistency, Support, and Entailment

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

For a hearer to be willing to update with a sentence, the update should not lead to the absurd state. And if a speaker is to assert a sentence correctly, it should not constitute a 'real' update in her information state.

Definitie 11 Let s be an information state.

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

With respect to the rare states of total information about the world, the notions of being allowed and supported coincide, and could be equated with truth, for non-modal statements that is. Likewise, in such states being forbidden amounts to falsity.

According to semantic intuition, a discourse is unacceptable if there is not at least *some* state that allows it. And if a sentence is not supported by *any* non-absurd state, which means that no speaker could sincerely utter it, then that sentence is judged unacceptable.

Definitie 12

1. ϕ is *consistent* iff there is some information state which allows ϕ
2. ϕ *is coherent* iff there is some non-absurd state that supports ϕ

Note that coherence implies consistency. Concerning the acceptability of a single sentence, it would suffice to require coherence. Still, it makes sense to distinguish both notions. A discourse may consist of a sequence of sentences, possibly uttered by different speakers in different information states. The acceptability of a discourse minimally requires that one by one the sentences are coherent. That does not imply that the discourse as a whole can be supported by a single information state. Hence, it does not imply that the discourse as a whole is consistent. The latter is an independent requirement for the acceptability of a discourse.

Since discourse consistency and sentence coherence are necessary conditions for acceptability, these semantic properties present us with criteria for testing the adequacy of a proposed semantics.

For this purpose, the notion of entailment is just as important. Entailment is not defined in the usual way in terms of truth, but in terms of sequential update and support:

Definitie 13 $\phi_1, \ldots, \phi_n \models \psi$ iff for all information states s such that $s[\phi_1] \ldots [\phi_n][\psi]$ exists, it holds that $s[\phi_1] \ldots [\phi_n]$ supports ψ

Below, some of the properties of the entailment relation are illustrated.

3.2 Equivalence

A suitable notion of equivalence may be expected to tell when two expressions can be substituted for each other in a meaning preserving way. Within update semantics, meaning is preserved if the update effects are. This being so, the usual definition of equivalence in terms of mutual entailment cannot be used. For example, $\exists x Px$ and $\exists y Py$ mutually entail each other, and so do $\exists x Px$ and Px, but, obviously, they cannot be replaced for each other in all contexts preserving update effects.

At the same time, it will also not do to require that ϕ and ψ are equivalent if they have exactly the same update effects. Under such a definition, $\exists x \exists y Rxy$ and $\exists y \exists x Rxy$ would not come out equivalent, and neither would $\exists x Px$ and $\exists x Px \wedge \exists x Px$. The reason for this is that the referent system of an information state not just keeps track of which variables and pegs are present, but also of the order in which they were introduced. Furthermore, there can be pegs around that are no longer associated with a variable. In view of this, the resulting information states are not required to be the *same*, but to be *similar*, where the notion of similarity ignores these less important differences between information states.

Definitie 14 Let $i, i' \in I, i = \langle r, g, w \rangle, i' = \langle r', g', w' \rangle$, with v and v' the domain of r and r', respectively; and let $s, s' \in S$.

1. *i is similar to i'* iff $v = v'$, $w = w'$, and $\forall x \in v\colon g(r(x)) = g'(r'(x))$
2. *s is similar to s'* iff $\forall i \in s\colon \exists i' \in s'\colon i$ is similar to i' and $\forall i' \in s'\colon \exists i \in s\colon i'$ is similar to i

Similarity is an equivalence relation.

Definitie 15 $\phi \equiv \psi$ iff for all information states s: $s[\phi]$ is similar to $s[\psi]$.

4 Illustrations

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

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

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

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

Feit 3 Whereas $\Diamond p \wedge \neg p$ is consistent, $\neg p \wedge \Diamond p$ is inconsistent.

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

Feit 4 Although consistent, $\Diamond p \wedge \neg p$ is incoherent.

An utterance of a sentence is incoherent if no *single* information state can support it.

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

Feit 5 $\Diamond p \wedge \neg p \not\models \Diamond p \wedge \neg p$

The reason behind all this is the *non-persistence* of formulae of the form $\Diamond \phi$: A state s may support $\Diamond \phi$, whereas a more informative state s' may be inconsistent with it. The non-persistence of modal formulae causes non-monotonicity of entailment:

Feit 6 Although $\Diamond p \models \Diamond p$, we have that $\Diamond p, \neg p \not\models \Diamond p$

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

4.1 Coreference and Modality

It is a characteristic feature of dynamic semantics that existential quantifiers can bind variables outside their scope. The variable in the second conjunct of (3) is bound by the quantifier in the first conjunct:

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

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

Exactly the same happens when s is updated with $\exists x(Px \wedge Qx)$:

Feit 7 $\exists x Px \wedge Qx \equiv \exists x(Px \wedge Qx)$

With the aid of the extended binding power of the existential quantifier a compositional and incremental account of cross-sentential anaphora can be given, and the same holds for donkey-anaphora.

Feit 8 $\exists x Px \rightarrow Qx \equiv \forall x(Px \rightarrow Qx)$

Equivalences such as these are characteristic of dynamic predicate logic.

Modal operators are transparent to the extended binding force of existential quantifiers. In (4), the occurrence of the variable within the scope of the *might*-operator is bound by the quantifier in the first conjunct:

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

In this case, the second conjunct only tests whether in the state that results after updating with the first conjunct there is at least one possibility

$\langle r[x/n], g[x/d], w\rangle$ such that $d \in w(Q)$. In particular, this means that among the possible values of x after updating with the whole sequence, there may be objects d that in no w have the property Q.

As is to be expected, both (5) and (6) are inconsistent:

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

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

But the following formula is *not* inconsistent:

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

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

However, unlike (7), (8) *is* inconsistent:

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

The brackets make a difference. In updating a state s with (8), some object d is chosen, and $s[x/d][Px \wedge \forall y\Diamond\neg Py]$ is performed. In all possibilities that remain after updating $s[x/d]$ with Px, d has the property P. But then $\forall y\Diamond\neg Py$ will be inconsistent with $s[x/d][Px]$. And this holds for each choice of d. Hence (8) is inconsistent.

This means that dynamic modal predicate logic lacks some features which characterize dynamic predicate logic. This point may be elaborated.

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

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

(9) There is someone hiding in the closet. He might be guilty.

$\exists x Qx \wedge \Diamond Px$

But the information state of your spouse would not support:

(10) There is someone hiding in the closet who might be guilty.

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

However, if the noise is a high-pitched voice, things are different. Now, your spouse knows it can't be your eldest, he already has a frog in his throat. In that case your spouse *can* say (10).

This also means that if your spouse yells (10) from upstairs, you can stay were you are, but if it is (9), you might run upstairs to check whether it is perhaps your aid that is hiding there.

So, there is a difference between (9) and (10),[3] and the semantics accounts for it:

Feit 9 $\exists x Px \wedge \Diamond Qx \not\equiv \exists x(Px \wedge \Diamond Qx)$

A similar observation applies to the following pair of examples.

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

(12) Anyone who is hiding in the closet might be guilty.
$\forall x(Qx \rightarrow \Diamond Px)$

Take the same situation again. Only in case your spouse heard some high-pitched voice, (12) is a correct utterance. In the other case, (12) is not supported by the information state of your spouse, and only (11) is left.

Feit 10 $\exists x Qx \rightarrow \Diamond Px \not\equiv \forall x(Qx \rightarrow \Diamond Px)$

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

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

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

---

3. And we thank David Beaver for pointing this out to us.

Consider the world that results when the the operations and relations mentioned in (13) are given their standard interpretation in the domain of real numbers. In that case (13) will be supported by any state consisting of possibilities in which only this world figures.

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

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

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

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

Updating with $\exists x\Diamond Px$ outputs as possible values of x only those d such that in some w compatible with the information in s, d has the property P in w. If $\Diamond Px$ is within the scope of $\exists x$, the consistency test is performed one by one for each $d \in D$, and those d are eliminated as possible values for x for which the test fails.[4]

5 Identity and Identification

Consider the following example:

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

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

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

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

The sequence of sentences (14) is coherent, and hence consistent. If it is continued with (15), everything remains consistent. But viewed as a single utterance, (14) followed by (15) would be incoherent.

There are several situations in which (14) can be coherently asserted. One is the situation in which the speaker is acquainted with the person who did it, but does not know his name — his name might be Alfred, his name might not be Alfred. However, also the opposite case, in which the speaker

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

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

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

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

Which is consistent and coherent, as you probably would like it to be.

5.1 Identification and Identifiers

The following two sentences are consistent:

(17) $\exists! x Px \wedge \forall y \Diamond(x = y) \wedge \forall y \Diamond(x \neq y)$

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

These sentences express ultimate forms of non-identification. If an information state supports (17), it is known that just one object has the property P, but not which object it is. If an information state supports (18), it is not known of which object a is the name.

Sometimes more information is available.

Definitie 16 Let α be a term, $s \in S$.

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

If a term α is an identifier in s, then s contains the information who α is (in at least some sense of *knowing who*). If α is not an identifier in s, then there is at least some doubt about who α is.

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

Feit 11

1. α is an identifier 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 $\models \exists x \Diamond x = \alpha \wedge \forall y (\Diamond y = \alpha \rightarrow y = x)$

Identifiers are epistemically rigid designators:

Feit 12 Let α and β be identifiers.

1. $\models \Diamond(\alpha = \beta) \rightarrow (\alpha = \beta)$
2. $\models (\alpha = \beta) \rightarrow \Box(\alpha = \beta)$

5.2 Why Identifiers are Needed

Identifiers are needed. Otherwise, if we are ignorant at the start, we can never really find out who is who, in the sense of coming to know the names of the objects we are talking about. The following definition and observation explain why this is so.

Definitie 17 Let $\langle r, g, w \rangle \in I, \langle r, g', w' \rangle \in I$.
$\langle r, g, w \rangle \simeq \langle r, g', w' \rangle$ iff there exists a bijection f from D onto D such that:

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 $\langle f(d_1), \ldots, f(d_n) \rangle \in w'(P)$

Feit 13 Let $\mathbf{0}$ be the minimal information state.
If $i \in \mathbf{0}[\phi_1] \ldots [\phi_n]$, then for every $i' \simeq i$, $i' \in \mathbf{0}[\phi_1] \ldots [\phi_n]$.

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

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

Definitie 18

1. Let $d \in D$. Then $this_d$ is a term
2. Let $i \in I$. Then $i(this_d) = d$

By definition, demonstratives are identifiers. Once they are added to the language, observation 13 made above, no longer holds. Expressions such as $this_d = a$ are now available, which can tell us which object a refers to.

Identifiers have a special logical role. Suppose the domain consists of two distinct individuals d and d'. We update the state of total ignorance with the following sentence.

(19) $(a \neq b)$

The resulting information state, s, supports

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

But s does not support:

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

Actually, s does not even allow (21), despite the fact that s supports the two instantiations with a and b — even though these are the names of all the objects around!

The state s does support (22) and (23):

(22) $\Diamond(this_d = a) \wedge \Diamond(this_{d'} = a)$
(23) $\forall x \Diamond(x = a)$

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

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

which can be straightforwardly derived from (23) by universal instantiation — or so it would seem. In other words, universal instantiation is not always valid. In particular, things may go wrong if one instantiates with a term which is not known to be an identifier. Likewise, existential generalization sometimes fails:

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

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

6 Prospects

We end by indicating some extensions of the system, restricting ourselves to issues which are immediately relevant to the topic of modality and coreference.

In the present paper, the assumption is made that the language users know which objects constitute the domain of discourse. Lifting this fixed domain assumption is relatively easy. The only important change that has to be made, is to keep track of the different stages of the ongoing interpretation process in an explicit manner. This is needed to enable accommodation of existential presuppositions of names and demonstratives in the proper way.[5] In most cases, presuppositions are not accommodated locally, but globally.

5. For example, along the lines of Zeevat 1992.

For the same reason, this additional structure is needed if the language is extended with anaphoric and non-anaphoric definite descriptions. Accessibility of the various stages of the interpretation process is also indispensable for an account of modal subordination.[6] Subordinated states should remain accessible, since subsequent sentences in the discourse may relate to those, rather than to the global, superordinate level. An analysis of this and other forms of subordination makes it possible to account for the fact that, under certain circumstances, quantifiers inside the scope of negation and the *might*-operator can bind occurrences of variables outside.

In the present paper, only epistemic modalities haven been investigated. Although the information based nature of dynamic semantics may suggest otherwise, this is not a principled limitation. Alethic modalities can be added, making it possible to implement the Kripkean distinction between metaphysical and epistemic necessity. For this purpose a set of metaphysically possible worlds is added to each possibility. Different possibilities may contain different alternative sets of such worlds. In this way, one can account for the learnability of what is metaphysically possible, necessary, and impossible.

Bibliography

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

Groenendijk, J. and Stokhof, M.: 1990, 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

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

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

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

6. See Roberts 1989.

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

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

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