CHAPTER NINE

Partial Information, Modality, And Intentionality

Fred Landman*

It is almost a truism to say that a semantic theory for natural language, which relates linguistic expressions with nonlinguistic interpretations, has to take into account the fact that language users are in general only partially informed about the intended domain of interpretation. Several semantic phenomena can only be adequately dealt with within a theory that takes this partiality seriously; vice versa, the semantics equally constrains the theory of partial information.

In this paper I will discuss some issues concerning the form that a semantic theory of information should take. Several of the issues I will bring up here, I have discussed in previous work (especially, Landman 1986), but I will follow a different line of presentation here and put some accents differently.

The main topic of the present paper will be the semantics of epistemic modals and identity statements. I will take as my starting point Hans Kamp's theory of Discourse Representation Theory (Kamp 1981), discuss its status as a theory of information, and point out some (obvious) problems it may come into if it is to deal with epistemic modals and their interplay with identity statements. The discussion of how to remedy such problems will lead me to some general discussion of partial semantics, and a sketch of a particular proposal for a partial semantics for Discourse Representation Theory for the phenomena in question. The semantics I will present takes a rather nonstandard view on what objects are. I will spend some time discussing how the semantics should be interpreted and how it fits in a broader theory of conversation and intensionality.

DISCOURSE REPRESENTATION THEORY

Discourse Representation Theory (DRT) is a rather fruitful semantic theory, developed by Hans Kamp (Kamp 1981) and Irene Heim (Heim 1982), for representing information that can be extracted out of sentences uttered in discourse. In Kamp's version, DRT contains a level of discourse representation structures at which for instance anaphoric links between discourse anaphora and their antecedents are represented, i.e., the links between the underlined expressions in (1) and (2):

(1) *A man* comes in. *He* wears a hat.
(2) If *a man* comes in, *he* wears a hat.

These structures get a semantic interpretation through recursive embedding conditions in a model. So the structure of the theory is:

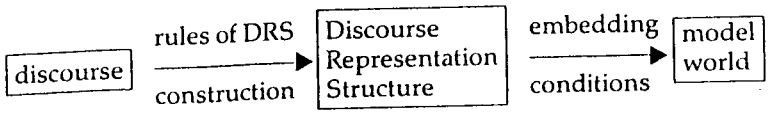

For a sentence like (1) a drs like D_1 is constructed:

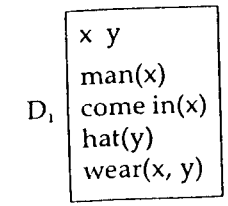

x and y are discourse referents; statements like man(x) are called discourse conditions. The crux of the analysis of discourse anaphora is that indefinites like *a man*, *a hat* do not introduce existential quantifiers but discourse referents, **variables**, that are bound by explicit or implicit quantifiers in the discourse. Discourse anaphora like *he* in (1) can be interpreted coreferentially with antecedents like *a man* only if they are in the scope of the quantifier that binds the antecedent discourse referent (variable). The embedding conditions for D_1 give the drs its existential force: D_1 is true in a model M if **there is** a truthful embedding for D_1 in M, a function g assigning to x and y objects in M such that in M g(x) is a man that comes in, g(y) is a hat and g(x) wears g(y).

Here the two sentences of the discourse are interpreted in the same box. Other constructions, like conditionals, universal expressions, negation, etc. introduce subordinated structure in the drs. (2) roughly has the structure D_2:

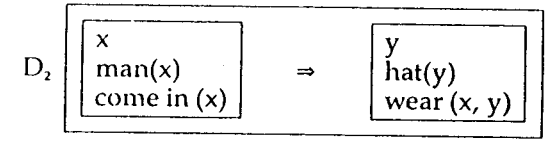

D_2 is true if every truthful embedding in M for the antecedent can be extended to a truthful embedding for the consequent in M (that is, every function g, assigning to x in M an object that is a man who comes in, can be extended to a function assigning to y a hat that g(x) wears. This is true in the model if for every man who comes in you can find a hat that he wears).

The interpretation of the conditional will give the sentence its universal force; moreover, *he* can be coreferential with *a man*, because (roughly) both are bound by the universal quantification linking the antecedent drs with the consequent drs. Note that we could not felicitously continue (2) as in (3):

(3) If *a man* comes in *he* wears a hat. **He* sits down.

The reason is that the *he* in the second sentence is not in the scope of the quantification that binds the discourse referent introduced in the antecedent of the conditional.

The drs, then, is a level of representation of information that is semantically interpreted through embedding conditions. The basic idea is that language users make a partial representation of the world, a "picture", on the basis of the sentences in the discourse that they hear and accept as true, and they try to match this picture with the world.

This interpretation of the theory brings out an important aspect: DRT is a theory of information with a **classical semantics**: it is based on a recursive specification of truth conditions (relative to a model, world) and hence based on an objective notion of truth. To use a metaphor, the semantics works like a telescope: in order to evaluate a discourse (representation structure) we point the telescope at the world and see whether we find the things that the discourse expresses in the world, we see whether it is true.

PARTIALITY

If DRT is to be used not only as a semantic theory, but also as a framework of representing information that language users do or do not have, it obviously has to have an opinion on partial information, on lack of information. That is, we have to be able to express in the theory what it means for a language user not to be informed about

whether something is true or not. How does DRT deal with this?

Suppose it is undetermined for our language user whether John has property P. Since the semantics is classical, in the model John either has P or he doesn't. So it is not at that level that the lack of information can be expressed. Here is the crucial observation. If the discourse representation stucture is empty, then every embedding function will verify it in the world M. As more information is entered into the drs, less and less embedding functions will be able to verify it in M. In view of this, we can regard discourse representation structures as **restrictions** on embedding functions. Adding conditions to the drs will eliminate previously available truthful embeddings, and only by adding conditions to the drs can embeddings be eliminated.

This means that the fact that it is not yet determined whether John has P or not, can only mean that there is no condition **in the drs** concerning P(j): **neither** P(j) nor ¬P(j) are represented in the drs. The world, the model, has of course an opinion on these, but **as far as the DRS is concerned**, no opinion is expressed on P(j) or ¬P(j). What is not mentioned in the drs is not yet decided as far as the language user is concerned. The language user does not have an opinion on P(j) if the drs does not contain either P(j) or ¬P(j).

But a problem arises. Such an analysis of partial information may be plausible for the case that I discussed, but it comes into problems if the drs contains expressions whose semantic function it is to **express** lack of information, partial information. For such expressions, the theory of information does not describe, so to say, from the outside what information is or is not available, but enters into the semantics for the drses themselves, into the embedding conditions. Expressions that function in this way are for instance epistemic modals *must* and *may* and conditionals.

An example may help to make clear that this is a problem for the simple perspective I have sketched. Take a sentence like (4):

(4) Maybe John is coming in, maybe he isn't.

This sentence **expresses** that, for a language user who utters or accepts it, it is undetermined whether John comes in. The drs is something like D_3:

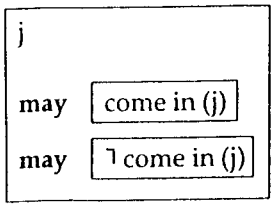

The question to be asked is: given that D_3 expresses the language user's lack of information concerning *John is coming in* (C(j)), and given the perspective on partial information discussed above, what should the embedding conditions for this drs be? In other words, what semantics should we give to **may**?

In the case of D_3, what **may** C(j) expresses is something like the following: even though I may not at the moment have the information that John is coming in, it is compatible with my present information that he is, it is possible that at a certain point I will have that information. Given the perspective on partial information that I mentioned, this 'at a certain point I will' has to mean that at a certain point C(j) will be explicitly in the drs (because that is, on this perspective, what it means to have that information).

So let us consider as a first guess (or strawman) the following semantics for **may**:

> **may** C(j) is true relative to g, a truthful embedding for D_3, if it is possible to extend D_3 to some D that explicitly contains C(j).

This analysis, of course, should be rejected out of hand. The problem is that it is an uninterpreted procedure that in no way captures the meaning of **may**. Any drs D_3 can be extended to some drs D containing C(j), in fact, it can be extended to a D' containing both C(j) and ¬C(j) (of course, no embedding function will verify it). With the above rule a **may**-statement cannot fail to be true.

So let us try to turn this condition into a more semantic one, like the following; our second guess/strawman:

> **may** C(j) is true relative to g, a truthful embedding for D_3 iff D_3 can possibly be extended to some D, containing explicitly C(j), such that g can be extended to a truthful embedding for D.

The difference between this condition and the previous one is that this is a semantic one. We now say: D_3 itself doesn't yet contain and verify C(j), but some possible extension of it does. The problem with it is that if in D_3 we make **may** C(j) true in this way, we cannot possibly make **may** ¬C(j) true in D_3 at the same time. This follows from the classical semantics. g, the embedding function for D_3 maps j onto some object such that some extension of g verifies C(j). Since no other discourse referents are introduced, that simply means that g verifies C(j) in M, and then, of course, no extension of g can verify ¬C(j) in M. In other words, this semantics fails to provide a consistent interpretation to D_3, D_3 cannot be true on this semantics.

Rather than trying to modify the clause once more, let us try to diagnose what is wrong with it.

PARTIAL SEMANTICS

Let me repeat the problematic clause for convenience.

(a) **may** C(j) is true relative to g, a truthful embedding for D_3 iff (b) D_3 can possibly be extended to some D, containing explicitly C(j), (c) such that g can be extended to a truthful embedding for D.

This clause was primarily an attempt to combine the notion of lack of information as being not yet represented in the drs with the idea that there is a truthful embedding for **may** C(j) in M if there can be a truthful embedding for C(j) in M. However, given the extensional semantics, we have to conclude that this *can* in the second clause is really misleading: it suggests that we are looking at possible extensions, it suggests a modality here, but the semantics does not provide this. Clearly we need a modal theory of some sort here. This is not the only problem, however. Suppose we do replace the last part (c) by a suitable modal theory (Like: (c') 'such that there is a truthful embedding of C(j) in some other world'). It will then be that modal theory, i.e., the structure of available possible worlds, and not whether something is or is not represented in the drs that will determine whether the language user has the information that C(j) or not. In other words, a correct semantics for epistemic modals will provide us with a **semantic** theory of information, distinct from the "syntactic," representational one introduced earlier, and it is this semantic notion of information and not the syntactic one that will play a role in our semantics: once we get the semantics of **may** right, the condition that the drs containing **may** C(j) can be extended to a drs containing explicitly C(j) (b) is totally superfluous. (In other words, the definition could be: (a) **may** C(j) is true relative to g, a truthful embedding for D_3 iff (c') g can be extended to a truthful embedding of C(j) in some other world.)

I have criticized the middle (b) and the end (c) of the above unfortunate clause, let's now turn to the beginning (a), that is the part 'g is a truthful embedding for **may** C(j) in M iff. . . .' I called the middle superfluous and the end wrong, the beginning I would characterize as nonsense.

Let me use the telescope metaphor once more. What we are doing, if we assume a semantics along the lines sketched is the following: we point the telescope at the world (M) and we ask: is it true in that world that maybe John is coming in? And this seems to be a nonsensical question to ask for modals like *maybe John is coming in*. In other words, it seems to be nonsense to try to embed *maybe John is coming in* into the world.

This is a serious point, because it affects not just a particular semantic clause but the whole semantic theory: apparently, when it comes to epistemic modals it won't do to take the notion "g is a truthful embedding for ϕ in M" as the basis of our semantic recursion.

Back to the telescope. Our telescope is a mechanism for inspecting semantic objects (like worlds). If we need a semantic notion of information in our theory then information states will be semantic objects, and hence our telescope can inspect them. So let us not point the telescope at the world, but at our information about the world. In other words, let us not try to match drses with the world, but with our information about the world. Now we again ask: is it true according to the information we are inspecting that John is coming in? The answer is yes if this information tells us so, no if it tells us that he doesn't, and undetermined if the information doesn't tell. Now ask: is it true according to the information that maybe John is coming in? This will be answered positively if the information can still possibly be extended to better information that tells us that John is coming in; negatively if such an extension is not possible.

Note that by changing the metaphor in this way, not only do the non-modal questions stay well formed, but now the modal questions make sense as well.

Let me here insert a warning about terminology. When I use the words "information (state)" and "true according to the information" here, I do so without intending to restrict myself to correct information, that is, information that is true of our world. Information states encode the things that language users accept in the course of conversation and there is no reason for not accepting non-factual information states, nor is there any reason why we cannot evaluate a sentence relative to non-factual information. Information states as I use them then correspond closely with Robert Stalnaker's common grounds and context sets (e.g., Stalnaker 1979), Frank Veltman's prejudice sets (Veltman 1976) and Angelika Kratzer's modal bases (for instance, Kratzer 1981).

The point I made above can be reformulated: modals are context dependent expressions: we do not ask whether **may** ϕ is true in the world, but only whether it is true given and in view of a contextually given information state. This context dependency of modals is well established in the literature (most strongly and systematically in Angelika Kratzer's work, see for instance, Kratzer 1981). So the very least we should do is extend the notion of context (already necessary for indexicals) with an information state (or a structure of information

states, if we want to be able to deal with different modalities) and assume that the semantic interpretation is sensitive to this context.

Context sensitivity and indexicality

Typical examples of context sensitive expressions are, of course, indexicals, like *I* and *now*. Such expressions are "rigid": regardless of whether they are embedded under modals, tenses, or attitude verbs, indexicals are interpreted relative to the outside context of use (*I* denotes the speaker in all of *I am ill, maybe I am ill, I will be ill, John thinks I am ill*). In order to interpret these sentences we do need to relativize the semantic recursion to the context (ϕ is true in w relative to c), but unlike the possible world parameter it does not play a real role in the semantic recursion. The context gives us the parameters that have to be set before we can start the semantic recursion. Or, as Kaplan 1979 and Stalnaker 1979 formulate it: the context determines what proposition the sentence expresses, given that, the semantic recursion on the possible worlds will determine whether it is true.

Not all context sensitivity is indexicality. Partee 1987 discusses the need to distinguish the general context of use from the local context. For instance, spatial adverbials like *nearby* are context dependent in that they need a spatial location to which they are related, but this spatial location need not come from the general context of use ('near to us') but can come from the local context: cf.

(5) Mary lives nearby
(6) John lives on State Street and Mary lives nearby.

The second sentence is ambigious in that it can mean 'near to us' or 'near to John.' Let us think of the local context of *Mary lives nearby* in the second sentence, as derived from the outside context and the interpretation of the first part of that sentence. Then indeed the interpretation of *nearby* is context sensitive, but not indexical: it would be indexical if we could only interpret it relative to the place of the outside context, but in fact it can also be interpreted relative to the local context.

Given that modals are context sensitive, let us ask whether they are indexical.

I will discuss this question in some detail here, because answering it is important for the discussion of modality later in this paper. My purpose there will be to argue in favour of one of the two plausible perspectives one can take on modals. Those two analyses are both context dependent, but not indexical. The arguments I will give there, however, are not arguments against an indexical analysis. Their

strength, then, would be undermined, if an indexical analysis is a viable alternative.

Concerning the indexical analysis, some of the discussion in Kratzer 1986 can be interpreted as a defense of this position, though it would be too strong to ascribe the claim that modals are nothing but indexicals to her there. The indexical position claims that just like *I am walking* can only be true in a world relative to a contextually provided speaker, *Maybe I am walking* can only be true relative to a contextually provided information state, contextually provided by the context of use.

This position differs drastically from most theories of conditionals and modals (including Kratzer 1981). Look at the following example:

(7) Maybe John killed the doctor. But if he was at home that night, he can't have killed her.

The *maybe* in the first sentence is interpreted relative to the contextually given information state, the sentence expresses that this information still allows for the possibility that John killed the doctor. If modals are indexical, then the *can't* in the second sentence (under the conditional) is interpreted relative to the same contextually given information state. But that would mean that the discourse can only be false or trivially true (if the antecedent of the conditional is false) with respect to that information state, because (8) is a contradiction:

(8) Maybe John killed her and he can't have killed her.

This example seems to suggest directly that modals are context sensitive, but not indexical, just like the mentioned spatial adverbials.

The only way we can get out of this problem is by assuming that all such problem cases, where we have a modal embedded in the consequent of a conditional, in fact have a **logical form** where the conditional is under the modal, more precisely *if ϕ then may ψ* has a logical form *may* (ϕ, ψ). This view fits, of course, very well with the idea of DRT (going back to Lewis 1975) that antecedents of conditionals only serve to **restrict** a quantifier (in this case *may*). This position is taken in Heim 1982 (for the examples she discusses there), where it is further assumed that this logical form is interpreted as: 'there is a possible world compatible with the contextual information where both ϕ and ψ are true.' This analysis fits with our tendency to read (9) as (10)

(9) If John comes to the party, Mary may come too
(10) Maybe John and Mary come to the party.

Heim's particular analysis is problematic for the following reason. Suppose I say to you:

(11) If I get my driver's licence next month, I may come and visit you.

What you don't know is that I know for sure that I won't have my driver's licence next month: I'm just being polite.

For sure, I did a bad thing. If you find out about it, you will be offended. Given my information, my statement was surely **incorrect**.

But was it false? I don't think so. I don't think that you can accuse me of **lying** to you (my response to such an accusation will be a treacherous, but nevertheless correct: "I wasn't lying, I said: **if** I get my licence...").

This makes sense if *may* is under the scope of the conditional, because then *maybe I get my driver's licence* is a quantity conversational implicature (rather than an entailment) of my statement (11), and violating a quantity implicature is not lying.

However, on the wide scope theory of **may** in Heim's version, my statement had the same meaning as (12):

(12) Maybe I get my driver's licence and I come and visit you.

This sentence can plausibly be argued to be a lie, because I know I won't get my driver's licence.

So Heim's analysis is not adequate. There is an alternative, however, that is more in line with Lewis' 1975 analysis of adverbs of quantification anyway: let us give the two-place **may** (ϕ,ψ) the following interpretation: 'of all the worlds allowed by the contextual information where ϕ holds, in some of those ψ holds.' We interpret a sentence where *may* seems to have only one argument, like *maybe John is coming*, as taking as implicit first argument the whole conversational information: **may** (S, John is coming), i.e., we take *Maybe John is coming* to mean: of all the possible worlds allowed by the context, in some of them John is coming. On the plausible pragmatic assumption that S is not empty, this will give us the correct interpretation.

We avoid now the problem with the example I gave, because the logical form **may** (ϕ,ψ) does not entail **may** (S,ϕ), but at most implicates that. In fact, what we are doing here is giving **may** (ϕ,ψ) an interpretation that is only subtly different from the interpretation that is in other theories given to $\phi \rightarrow$ **may** ψ.

The indexicality theory of modals then makes the following claim: whenever we see a modal that in surface form is embedded under other operations, like conditionals, in logical form it is not embedded and we don't find modals that we have to interpret **under the scope** of conditionals or other modals (because they don't have scope).

This simplifies our logical, semantical work considerably. For

instance, we don't have to worry anymore about the semantics of conditionals that are not first degree, like $\phi \rightarrow (\psi \rightarrow \kappa)$: we simply won't interpret this directly, but give it the logical form: $(\phi \wedge \psi) \rightarrow \kappa$ (i.e., treating ϕ and ψ as restricting the same main quantifier): basically, we work all modals and conditionals up to the highest possible level).

Now, because this is not an approach that is systematically worked out and defended, I cannot argue that it is inadequate. But let me express some serious worries.

The disturbing aspect of this approach is the degree to which it has to rely on logical forms that are more and more remote from the surface forms of the expressions. In the above example, the use of logical form was still relatively mild (just move the **may** over the conditional). But when we look at what would have to happen with sentences where, say, the consequent of the conditional is itself logically complex and has modals embedded in it, the prospects of finding a systematic mechanism that will assign the right logical forms to sentences become highly dubious to me. For instance, in the following example we would have to work the modal up out of the conjunction, without it having scope over the first conjunct, and without losing the fact that the *it* in the second conjunct has *a hat*, introduced in the first conjunct, as its antecedent (and hence that should be accessible):

(13) If a man comes in, he wears a hat and he may give it to you.
$\phi \rightarrow (\psi \wedge$ **may** $\kappa)$.

This means that we will have to give this sentence a logical form like: $(\phi \rightarrow \psi) \wedge$ **may** $(\phi \wedge \psi, \kappa)$. Now, it may be possible to find such a rule and find a rationale for it, but at present it seems rather hopeless to me.

A discourse representation theory that is able to interpret the **may** in situ in this example does not need any logical reconstruction of this sort of the discourse and seems to be more attractive (to me at least) for that reason.

But such a theory has to deal with the logical problem of iterations of logical operations and modals and this is a highly non-trivial task (this is a very hard problem that I cannot go into here, see Landman 1986 for extensive discussion).

It is very hard to give knockdown empirical arguments in favor of either one position (modals as indexicals, vs. modals as context sensitive expressions, that can be sensitive to the local context) for various reasons. One is the vagueness of the present state of the indexical theory (if one logical form makes the wrong predictions, you can always draw back to another one). Another is the subtlety and vagueness of the data: it is just very hard to find strong intuitions about

iterated modalities. Another point is the closeness of the approaches. For instance: the analysis of **may** (ϕ,ψ) that I sketched above differs only subtly from the analysis of $\phi\rightarrow$ **may** ψ that you would find on the other approach, basically because in the other theory, the distinction between $\phi\rightarrow$ **may** ψ, **may** $(\phi\rightarrow\psi)$ and **may** $(\phi\wedge\psi)$ is very slight. Finally, an approach that says that modals and conditionals can have scope with respect to each other, will tend to assume that sentences with embedded modals show scope ambiguities: the fact that the theory allows, say, the sentence 'if ϕ then may ψ' to mean either one of the above three almost equivalent formulas doesn't make the task of deciding what it means and what it doesn't mean easier.

Let me discuss two more examples that bear on the question of the scope of modals in conditionals.
Let us listen to the following story:

> "There was this boy in my hometown before the war. We were in love. I'm not sure whether I still want to marry him. I don't even know whether he is still alive. Moreover, I don't know whether I will ever go back. What's the point: maybe I'll go back and find out that he is dead. **But if I ever go back, I may marry him.**"

I tend to think that the last sentence, in the context set up, has a reading where it is false (to formulate it overly cautiously), a reading where it is incompatible with the one but last sentence. This is the narrow scope reading: the sentence is true if every possible situation (allowed by the context) where she goes back, still allows for her marrying him, but this is excluded in the possible situation where she goes back and finds out that he is dead. The indexical theory predicts that the only reading that the sentence has in this context is: "it is possible that I go back and marry him." Again, a scope theory does not exclude this reading, but tells you that it has the other as well.

A second example has to do with donkey-sentences. We hear the following description:

> "This guy is so paranoid. He sees Jack the Ripper everywhere. His methodology is: **If a prostitute is killed, then Jack may have killed her.**"

In this context, this sentence does not mean that out of all the possible cases where a prostitute is killed, in some (possible ones) Jack did it, but rather that for each possible case, the possibility that Jack killed her cannot be excluded. Again, the indexical theory seems to predict that this reading doesn't exist.

As I mentioned before, I won't in this paper go into the analysis of

iterations, in fact, I will for simplicity use an analysis that for iterations makes blatantly wrong predictions (but see Landman 1986 for how it can be modified). Yet I hope to have given some motivation for a central assumption that I will make, an assumption that I share with most of the work on conditionals (especially the work closest to the present paper, Kratzer 1977, 1981; Veltman 1976, 1981, 1985, 1986; Landman 1986), that there are important differences between the context dependency of indexicals and of modals.

The conclusion, then, that can be drawn from all this is that information states play a more fundamental semantic role in the analysis of modals than the context of use (i.e., whatever deals with indexicals). The information state crucially enters into the semantic recursion (to give us the local context), while the context of use, so to say, stays on the outside. Maybe it clarifies the discussion to point out the difference with a more standard modal logic. In a standard modal logic the semantics determines, given the context, the possible worlds in which a complex sentence is true, given the possible worlds in which its parts are true. On the present perspective the semantics determines, given the context, the information states in virtue of which the complex sentence is true, given the information states in virtue of which its parts are true. Theories that have the latter form are given in the mentioned works by Veltman, Kratzer and myself.

The theory in Kratzer 1981 is in fact a mixed theory, where the semantics determines the world-information state pairs in virtue of which a complex sentence is true, given the world-information state pairs where its parts are true. In her theory the world part is used for non-modal expressions and the information part for modals. The distinction is important and leads to subtle differences with the approach followed by Frank Veltman and me. Since the semantics for non-modal expressions is classical (a standard possible world semantics) the notion of a speaker having partial information with respect to a non-modal expression is not given by the semantics directly, but has to be defined using the notions that the semantic theory provides (this can easily be done in the classical way: the information state as a set of alternatives provides the information that ϕ if ϕ is true in all alternatives). Such a definition is an essential part of the theory, because epistemic uncertainty is not a semantic borderline phenomenon that can be isolated in the grammar by saying 'there are a few expressions (*must* and *may*) that are information sensitive, they are context dependent.' Since the modals and conditionals interact strongly with the other connectives (like negation and disjunction), epistemic uncertainty and the use of devices to express among others what information we have or do not have is at the heart of semantics (although we

might decide to call it pragmatics, rather than semantics). A more classical setup like Kratzer's 1981 theory makes predictions about the interplay between modal and non-modal expressions that are interestingly different from a theory that takes the notion of an information state as the basis for the semantic recursion for both modal and non-modal expressions. One such case (the relation between ϕ and **must** ϕ) I will come back to in the next section.

I said earlier, when I talked about the telescope metaphor, that it is sensible not just for modals but also for non-modal expressions to ask whether they are true or false relative to the information. Given what I just said about epistemic uncertainty being at the heart of semantics it becomes attractive to develop this metaphor (pointing at the information) into an alternative to the classical theory based on the classical metaphor (pointing at the world) (that is, the theory that brought us into problems to start with). This we could call a **partial semantics**: the semantics does not specify conditions for truth/falsity per se (though truth and falsity can be defined as the borderline case of total information), but **the semantics specifies recursive conditions for truth on the basis of the information/falsity on the basis of the information**. Again, the terminology is meant to be neutral, I don't care whether you want to call them truth-on-the-basis-of-the-information conditions, evidence conditions, verification conditions, or assertability conditions (although each choice of terminology has its own welcome and unwelcome associations; assertability (Dummett's term) is probably the most neutral), the important point is that (even though a particular semantic theory developed in this framework can be classical) the framework is non-classical in that it does not take the capacity to recognize whether a sentence is true or false in a situation to be at the heart of the semantic recursion, but the capacity to recognize whether an information state does or does not justify an assertion.

Before I continue to fill in some details, let me draw a general moral for discourse representation theory out of the preceding discussion. The work on modals shows that we need to recognize a semantic level of partial information. Besides the level of discourse syntax (drses), we need a level of discourse semantics as a level of partial information. If we incorporate this in the theory as a partial semantics, then the structure of the theory will be as follows:

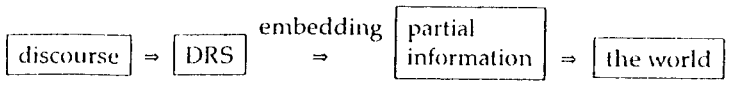

Some Aspects Of Data Semantics

In this section I will do two things. I will first give a short sketch of some relevant aspects of data semantics. Apart from some remarks on discourse anaphora, the points I will discuss here are not new. Most of them can already be found in Veltman 1981, one example (15-18) comes from Veltman 1987, some more discussion from Landman 1986. Since I have to be brief here, I have to refer the reader to the works by Frank Veltman and by me in the references for details, motivation, etc. After this, I will compare the data semantic analysis of modals with the more well known possible world style analysis.

The models for the theory are familiar Kripke models (with some conditions on them that I won't go into here).
A model is a tuple $<S, \leq, D, i^+, i^->$.
S is a set of possible information states, partially ordered by a relation of information extension $\leq$.
For every information state s, D(s) is the domain of s. I will assume that domains grow as information grows: i.e., if $s \leq s'$ then $D(s) \subseteq D(s')$.
i^+ and i^- assign to every predicate and information state a positive and negative extension respectively, i.e., $i^+(P, s)$ is the set of objects that have P on the basis of s, $i^-(P, s)$ is the set of objects that (definitely) do not have P on the basis of s. Obvious conditions are:
$i^+(P, s) \cap i^-(P, s) = \phi$ (the positive and negative extension of a predicate do not overlap)
$i^+(P, s) \cup i^-(P, s) \subseteq D(s)$ (both the positive and the negative extension of P are subsets of the domain)
if $s \leq s'$ then $i^+(P, s) \subseteq i^+(P, s')$ and $i^-(P, s) \subseteq i^-(P, s')$ (not only the domains, but also the interpretations grow)
Several other things should be made explicit (concerning for instance total information states), but I will omit such discussion here.

The semantics defines recursively the following two relations:

g $\Vdash_s$ C: embedding g verifies condition C **on the basis of** s
g $\dashv\Vdash_s$ C: embedding g falsifies condition C **on the basis of** s.

Again I should be much more explicit about the semantics, in particular about how this fits in with a full DRT-semantics. I hope to do that in another paper. The cases that I will discuss do not particularly depend on the suppressed parts.

Partial and total information

Given the models introduced above, the clauses for atomic conditions and their negations are unsurprising:

$g \Vdash\!\dashv_s P(x)$ iff $g(x) \in i^+(P, s)$
$g \dashv\!\Vdash_s P(x)$ iff $g(x) \in i^-(P, s)$
$g \Vdash\!\dashv_s \neg A$ iff $g \dashv\!\Vdash_s A$
$g \dashv\!\Vdash_s \neg A$ iff $g \Vdash\!\dashv_s A$

On partial information both P(x) and ¬P(x) can be undefined. On total information all epistemic uncertainty will be resolved, the logic becomes classical logic. So classical logic indeed is the limit case of total information.

Instability of may

$g \Vdash\!\dashv_s$ **may** A iff for some $s' \geq s$ and some $g_A \supseteq g$: $g_A \Vdash\!\dashv_{s'} A$.
(Here and in what follows I am assuming the formulation of DRT that is given in Landman 1987 as Theory 4.)

maybe John is coming in is unstable (non-monotonic) under information growth: it can be true on the basis of partial information (because there is still a possible extension where John is coming in), but false on the basis of better information (because the previous extension is no longer possible). This is different for atomic expressions: *John is coming in* is stable: once it is in the information, it will stay there (under information growth).

The unstability of **may** can be observed as well in (5):

(14) Maybe John killed the doctor. But if he was at home that night, he didn't kill her.

The contextual information state still allows an extension to an information state containing the possible fact that John killed the doctor, but the information states with respect to which the consequent of the conditional is evaluated (which are information states extending the contextual one with the antecedent information) cannot be extended to such an information state any more.

The indirectness of must

$g \Vdash\!\dashv_s$ **must** A iff for every $s' \geq s$ there is an $s'' \geq s'$ such that $g \Vdash\!\dashv_{s''} A$

(I am simplifying by ignoring the assignment extensions here.) This means: in whatever way the information may grow, in the end A comes in the information. The distinction between A and **must** A is the distinction between direct and indirect information: A is stronger than **must** A (at least for atomic A): A is directly true on the basis of s if A is one of the facts in s, it is indirectly true on the basis of s if there is no way of extending the information to an information state

where ¬A holds. (If A is stable and directly true it is also indirectly, but A can be indirectly true without being yet directly true.)

Of course, one may wonder: isn't the notion of direct information a philosophically unfounded notion? Isn't all information when it comes to it indirect, based on inference? And if so, why try to incorporate a notion of direct information in the semantics?

The answer to these questions, which I have defended at length in Landman 1986, relies on the distinction between (semantic) information growth and (pragmatic) information refinement. Notions like partial information, total information, and with them direct and indirect information, are strongly context dependent. Information is only total or direct with respect to a standard of precision. What is a simple piece of information under a loose standard of precision may turn out to be a highly complicated structure of facts under a more precise standard; what is a total description of the world may turn out to be highly partial if on a more fine-grained standard of precision other facts are taken into account; moreover, pragmatic accommodation can introduce possibilities that were previously not considered (for accommodation, see Lewis 1979; the fact that pragmatic accommodation can introduce possibilities is a crucial aspect of Stalnaker's conception of possible worlds (see Stalnaker 1984) and the way they allow for partiality).

So we need to allow for the possibility that pragmatic accommodation changes our initial information state and/or its relation to other information states. This is not incompatible with a semantic notion of total information or direct information. The sense in which all information is indirect is pragmatic: we can always refine our information pragmatically in such a way that what was direct before turns out to be indirect according to the new standard of precision. The situation is similar to the analysis of vagueness (Kamp 1975): here too we want to semantically relate a vague predicate to all possible ways of making it precise. But 'precise' means: precise with respect to a certain standard of precision; a predicate is inherently vague if we can always redraw the standard of precision is such a way that the sharp extensions according to the old standard, turn out to be vague with respect to the new standard (cf. also Lewis 1979's discussion of the classical example *France is hexagonal*).

But having recognized the strong influence of context and accommodation, one should also recognize the semantic need for notions like total information and the direct/indirect distinction. For instance, though we realize that we can accommodate a possibility that was previously not considered, therewith making **may** A true, we should at the same time realize that it involved **correction**, that after we have

accepted **may** A, we are no longer willing to accept **must** ¬A. This means that our semantic theory should capture the fact that, when evaluated with respect to one and the same information state **may** A and **must** ¬A are incompatible. And we can do that by distinguishing between information growth, the relation between an information state and its extensions, relative to a standard of precision (a background context), and information refinement, the relation between one system of information states and its refinements.

Similarly, there are semantic reasons for introducing a difference between direct and indirect information. That *John must be home* is a weaker statement than *John is home* has often been observed (for instance, Karttunen 1972) and the distinction is neatly shown in examples (15)-(18):

(15) I am hungry.
(16) ?I must be hungry.
(you can say (16), but it seems as if you infer it from external, indirect cues)

(17) You must be hungry.
(18) ?You are hungry.
(again, you can say (18), but it is normally impolite to inform someone about something which they normally have direct information about and you don't).

A semantic theory should represent these distinctions, and the proper way to do that, it seems to me, is to say that even though in the end the evidence on which my utterance of (15) is based may be indirect as well, still there is a clear sense in which, relative to the standard of precision of a normal context, (15) is direct in a way that (16) is not. This is what the data semantic clauses intend to capture. I will come back to the semantics of *must* later in this section.

Negation of conditionals

$g \Vdash_s A \rightarrow B$ iff for every $s' \geq s$ and every $g_A \supseteq g$ such that $g_A \Vdash_{s'} A$ there is an $s'' \geq s'$ and a $g_B \supseteq g_A$ such that $g_B \Vdash_{s''} B$.

The basic scheme of the conditional is: every extension verifying the antecedent can be extended to an extension verifying the consequent. In this paper I won't be concerned with the semantics of the conditional; for extensive discussion of this clause and its modifications and alternatives, see Landman 1986, Veltman 1985. I will make some

remarks about the corresponding falsifying clause, though. The obvious data semantic clause (given the above verification clause) is the following:

$g \dashv\Vdash_s A \rightarrow B$ iff for some $s' \geq s$ and some $g_A \supseteq g$: $g_A \Vdash_{s'} A$ and $g_A \dashv\Vdash_{s'} B$.

This means that $\neg(A \rightarrow B)$ is going to be (weakly) equivalent to **may** $(A \wedge \neg B)$

A lot has been written on the proper analysis of negations of conditional sentences. I here want to indicate that donkey sentences and discourse anaphora may add some new arguments in favour of the above analysis. Take the following examples:

(19) If a farmer owns a donkey, he beats it.
(20) It's not the case that if a farmer owns a donkey, he beats it.

Now look at (21)-(24):

(21) It's not the case that if *a farmer* owns a donkey, he beats it. **He* owns a cow as well.
(22) *A farmer* owns a donkey and doesn't beat it. *He* owns a cow as well.
(23) If a farmer owns a donkey, he doesn't beat it.
(24) Maybe *a farmer* owns a donkey and doesn't beat it. **He* owns a cow as well.

(19) is the classical donkey sentence, which quantifies universally over farmer-donkey pairs. What does (20) mean? Classical logic tells us that (20) is equivalent to the first sentence in (22) ($\neg(A \rightarrow B)$ iff $A \wedge \neg B$), so (21) and (22) should have the same meaning. But they don't: apart from the fact that (22) seems to make a much stronger claim than (21), the anaphora is not possible in (21), while it is in (22).

Stalnaker 1968 tells us that (20) is equivalent to (23) ($\neg(A \rightarrow B)$ iff $A \rightarrow \neg B$). Since (23) is a donkey sentence as well, it means: for every farmer-donkey pair such that the farmer owns the donkey, the farmer doesn't beat the donkey. Clearly this is not right for donkey sentences, (20) certainly does not mean (23). The data semantics analysis tells us that (20) is equivalent to the first sentence in (24) (and hence (21) and (24) are equivalent). Not only does this intuitively sound quite plausible, but it makes the right predication about anaphora possibilities as well.

This connection seems to be even strengthened if we look at modal subordination cases. Modal subordination (see Roberts 1987) is the phenomenon that under a certain continuity of modality anaphora is possible, but stays "in the scope of the modals," as in:

(25) *A tiger* may come in. *It* would eat you first.

We have observed above that in general anaphora outside a modal or a conditional, or the negation of a conditional cannot take an antecedent inside those. But they can take such an antecedent on the modal subordination interpretation. The observation that I want to make is that for all three of the following examples (26)-(28) a modal subordination continuation seems possible and on that reading they all seem to be equivalent to (29):

(26) Maybe *some farmer* doesn't own a donkey. For instance, maybe *he* owns a cow instead.
(27) It's not the case that *every farmer* owns a donkey. For instance, maybe *he* owns a cow instead.
(28) It's not the case that if *someone* is a farmer, he owns a donkey. For instance, maybe *he* owns a cow instead.
(29) Maybe some farmer doesn't own a donkey but a cow.

Again, I think that the parallel suggests that the data semantic analysis is on the right track.

Note that this analysis of the negation of conditionals will bring the instability of **may** into the logic of the conditionals. A consequence of this is that the principle of *modus tollens* does not hold generally (see Veltman 1981, 1985; Landman 1986).

Modus Tollens

Take the following argument: Maybe there will be war, maybe there won't be war. But if B wins the elections there must be war, war is unavoidable.
With modus tollens we can conclude:
may W, **may** ¬W (= ¬ **must** W), B → **must** W, hence ¬B. In other words, if modus tollens is valid, the above argument implies: B doesn't win the elections. This is of course much too strong. In data logic we get the following dichotomy:

If ¬Y is stable, we have indeed: X → Y, ¬Y, hence ¬X
If ¬Y is instable, we only have: X → Y, ¬Y, hence **may** ¬X

Since ¬ **must** A is instable, the above argument only leads to the weaker (and correct) conclusion: *maybe B doesn't win*.

Indirectness of must *in data semantics and possible world semantics*

In the next section, I will argue that the data semantics analysis of modals which I have sketched here forces us to adopt a radically non-

standard analysis of identity statements. Although the above analysis differs on all the points mentioned from what we are used to in possible world semantics, we may wonder whether a rich framework like possible world semantics is not able to take over the desirable aspects of the data semantic analysis. This is an important question, because, if we can incorporate the important aspects of data semantics in a possible world-style theory, it is quite likely that we will be able to combine the analysis of modals with a standard, classical analysis of identity statements. Thus, the argument in the next section, that the analysis of modality unavoidably leads to a non-classical theory of identity, would be unfounded.

I will argue here that the distinctions on which data semantics is based — in particular the distinction between direct and indirect information — cannot be captured correctly in standard possible world semantics (of course it can be captured in a less standard possible world semantics: data semantics is a less standard possible world semantics); moreover I will defend the data semantics analysis of *must* against the analyses that are possible in possible world semantics.

Let us see how possible world semantics could take over a data semantic analysis.

In possible world semantics, sentences denote propositions, sets of possible worlds. Following Stalnaker's work (e.g., Stalnaker 1979), there are two notions of information state. The first is what Stalnaker calls the common ground and what Kratzer calls the modal base, a set of propositions which are regarded as common information by the speech participants. The second is the context set, a set of alternatives (i.e., possible worlds). This is the set of all worlds compatible with the information in the common ground, i.e., the set of all worlds in which every proposition in the common ground is true (technically, this is equivalent to the closure of the common ground under logical consequence and conjunction).

Let us ask the following question: how can we capture the distinction between direct and indirect information, i.e., between A and **must** A in such a theory?

Already in the oldest work on *must* and *may* (Karttunen 1972) it is assumed that the distinction can be analysed as that between what is in the information and what is not yet in the information, but follows from it. Let C be a common ground and let ∩C, the intersection of all the propositions in C (i.e., the set of worlds in which every proposition in C is true) be its corresponding context set. Then we can define: A is true relative to common ground C if A ∈ C; *must* A is true relative to common ground C if A is true in every world in ∩C.

So indeed A is true relative to C if A is in C; *must* A is true relative to C if A follows from C.

The common ground is not closed under logical consequence, the context set is. So something can follow from the context set, without being explicitly (directly) in the common ground. In that situation it would be indirectly true.

However, as I argue in Landman 1986, this leads to a notion of direct information that is devoid of any content. Which propositions can provide direct evidence for the truth of A? In the present analysis, hardly anything but A itself gives us sufficient information to conclude that A is directly true; any other information gives us at most **must** A (the reason is that C is not closed under logical consequence). To give an example, not even if we have the information that John is sixteen years old can we conclude that "John is older than three" is directly true (because the first proposition can be in C without the second being in C). So we can introduce a direct/indirect distinction in this way, but it gives us a useless notion of direct information.

So we cannot use the logical gap between the common ground and the context set to characterize the direct/indirect distinction. In this framework, then, the only thing we can possibly try to use to define the direct/indirect contrast in a non-void way, is the fact that the corresponding context set, in virtue of its being a set of alternatives, is a partial information state. With this I mean the following: the context set is the set of all worlds compatible with the information. This set is constrained by the context, so it doesn't contain all possible worlds, but since our information is not total, it doesn't eliminate all possible worlds (except one) either: different alternatives for what the facts in fact are still open in the context set. In this sense, the notion of a context set comes very close to the data semantic notion of an information state. So the option that we can try to work out is to take the data semantic clauses for A and *must* A and impose them on the notion of context set. For the truth conditions of A, we have no choice but to follow the standard possible world practice: A is true relative to $\cap$C iff $\forall w \in \cap C$: A is true relative to w. For *must* A, we take the data semantics clause: *must* A is true relative to $\cap$C iff for every context set X extending $\cap$C, there is a context set Y, extending X such that A is true relative to Y.

What does it mean that X extends $\cap$C? It means that X contains better information than $\cap$C, which in possible world semantics means that X has eliminated certain alternatives that $\cap$C hasn't, i.e. X is a subset of $\cap$C. Given this, it is easy to see that these two notions of direct and indirect information collapse: A is true relative to $\cap$C iff *must* A is true relative to $\cap$C.

This exhausts the possibilities of the framework as it is, so we see that in the framework as I have presented it above, there is no sensible way of drawing the distinction between direct and indirect information.

In the possible world framework, what I think is the only way out is the way that Kratzer 1981 takes.

Modals, in her theory, are not just interpreted relative to one information state, the modal base, but relative to two: the modal base and the ordering source. The modal base C (as before a set of propositions) determines what kind of modality is involved (epistemic, deontic, etc.). The ordering source O, which is also a set of propositions, orders the worlds in $\cap$C. Leaving aside certain complications that are irrelevant for our purposes, O determines within $\cap$C a subset $\cap$O.

Here is an example of the usefulness of this notion. Suppose C is a deontic modal base. Then $\cap$C is the set of all deontically possible worlds. C determines the truth of sentences like: *I must not steal*: in view of what C allows and forbids, I must not steal, and this means that the proposition that I steal is incompatible with C. However, certain modals seem to make a less categorical statement than the previous one, like *I should give to the poor*. It's just too strong to claim that *I don't give to the poor* is incompatible with C: the above modal statement is more an 'adhortative' than an 'imperative.' For this reason, Kratzer assumes that within the set of deontically possible worlds, the ordering source selects out a subset of which we can think as those deontically possible worlds in which the deontic ideal that the ordering source contains (like 'the ideal of a good life') is realized. *I should give to the poor* is true relative to C and O if in all those permissible worlds (worlds in C) which are worlds in which I realize the ideals of a good life, over and above the fact that I don't do anything that isn't forbidden, if in all those worlds (and they are the worlds in O) I give to the poor.

I have no problem with the notion of ordering source as such. In fact, I think Kratzer makes a compelling case for them in her paper. I will here only be concerned with the question whether they are the right instrument to get us out of our problems.

This is the idea:
we keep the definition of A being directly true relative to C if A is true in every world in $\cap$C. We now redefine 'A is indirectly true (or **must** A is true) relative to C' as: A is true in every world in $\cap$O. Presumably we should think of $\cap$C as the set of epistemic alternatives and $\cap$O as the set of epistemically trustworthy alternatives.

In this theory A is stronger than **must** A (if A is true throughout $\cap$C, it is certainly true in $\cap$O), and also, **must** A can be true, without

A being yet true (relative to C and O). Hence, we have been able to create a distinction in possible world semantics after all.

However, there is a crucial difference between this theory, and the data semantic one I introduced above.

Let us consider a situation where we are willing to assert **must** A, but not yet A. For instance, you and I are walking over Cornell campus, you're figuring your way out with a map and finally you say: 'This must be Morrill Hall.' I am in a position to confirm this, I know the building, so I say: 'Indeed, this is Morrill Hall.' Let us assume that your information state C before my utterance was one where **must** A is true, but not yet A. The situation according to data semantics is that A is true in every total extension of C, but not yet in C itself. The situation according to Kratzer is that A is true in every world in $\cap O$, but not in every world in $\cap C$, which means that in some world in $\cap C$ A is false: in other words, there is an epistemic alternative, but not one very close to the ideal O, in which A is false. In Kratzer's terminology: someone who accepts **must** A in this situation allows for a slight possibility that A is false (in fact, there is a modal operator which we can call SP (for slight possibility) that expresses this: SP(A) is true relative to C and O if there is a world in $\cap C - \cap O$ where A is true).

Both theories claim the **must** A is accepted on the basis of indirect evidence. The difference between the two theories lies in what they think indirect evidence is. Data semantics claims that a language user, although she distinguishes between direct and indirect evidence, still treats indirect evidence as **good** evidence: someone who accepts **must** A does not leave open the possibility that $\neg A$ any more than someone who accepts A does. Kratzer's theory, on the other hand claims that a language user treats indirect evidence as **bad** evidence (or relatively bad evidence): it is not possible to have indirect evidence that A without leaving open at least the slight possibility that $\neg A$, while direct evidence that A is inconsistent with even the slightest possibility that $\neg A$.

We observe that the theories make different predictions. Kratzer predicts that in the situation described in fact **must** A is equivalent to SP($\neg A$) $\wedge$ **must** A. Data semantics predicts that SP($\neg A$) and **must** A should be inconsistent. Now look at the following sentence:

(30) There is a slight possibility that this is not Morrill Hall, but it must be Morrill Hall.

Kratzer predicts that this sentence should be perfectly all right, and in fact, in the situation described equivalent to (31)

(31) This must be Morrill Hall.

To me, this seems to be wrong: (30) is not equivalent to (31), but feels as inconsistent as (32):

(32) There is a slight possibility that this is not Morrill Hall, but it is Morrill Hall.

We should be careful, though. What about (33):

(33) Although it is logically possible that John is not the murderer, all the evidence we have gathered tells us that he must be the murderer.

This example (which I owe to Angelika Kratzer) does not seem to be inconsistent and seems to have the structure SP($\neg A$) $\wedge$ **must** A. However, I don't think that this is a convincing example. The second sentence (the **must**) is evaluated relative to an **epistemic** modal base. But the first modal is not related to the **same** epistemic modal base, but to a (less informative) **logical** modal base (specifying what is logically, rather than epistemically possible). Of course, the data semantics' claim is only that SP($\neg A$) and **must** A are incompatible if related to the same modal base, so this example is not a counterexample to it. On the contrary, I would even say, we can take the fact that we make explicit that in the first sentence a different modality is involved than in the second as an indication that we want to avoid a contradictory interpretation as in (30).

However, what should we say about the following example? You have guided me to Morrill Hall by use of the map and you say:

(34) This must be Morrill Hall. Well, of course there's a possibility that I made a mistake, that the map is inaccurate, and that this isn't Morrill Hall.

This also doesn't seem to be inconsistent in the way (30) is, and it doesn't necessarily have to involve the notion of **logical** possibility. So this would be a more convincing example in favour of Kratzer's theory. Another example that seems perfect I owe to one of the anonymous referees:

(35) John must not have heard what I said, but if he did, he can tell you.

However, I don't think that these examples are arguments for Kratzer's theory.

Let us look at (35) first. The case is slightly different from the other case, because the second sentence, if evaluated relative to the modal base relative to which the first is true would not be a contradiction but a tautology. However, the *but* indicates a contrast and what it does in

this context is indicate that a slightly different modal base may be relevant. That this is so is shown by example (36):

(36) John **didn't** hear what I said, but if he did, he can tell you.

Here the first sentence is not a modal at all, yet the sentence is perfect. If so, then we have every reason to assume that whatever makes (36) well-formed makes (35) well-formed as well.

Concerning (34), I have the strong impression that what is going on here is accommodation. The speaker makes a strong statement: this must be Morrill Hall. Then she wants to be more cautious, she **weakens** her statement slightly by bringing certain possibilities into the discourse. Consider (37):

(37) This must be Morrill Hall. Well, there is a slight possibility that it isn't Morrill Hall. But still, this must be Morrill Hall.

Unlike (34), (37) sounds to me as contradictory as (30). The explanation that I gave for (34) and a data semantics explain this. The second sentence would be inconsistent with the first if interpreted on the same information state. So instead the information state is (pragmatically) changed: a possibility is accommodated. But as soon as the second sentence is accepted as true, the first is no longer acceptable, because they are inconsistent. So adding the first sentence after the second has been accommodated results in a contradiction indeed.

On Kratzer's theory, the only thing that is wrong with (37) is that it is repetitive: given the first sentence, the second and the third do not add any new information. To me, that seems to be inadequate.

Summarizing: (37) shows that (34) is a case of accommodation or correction, that in the second part of (34) the claim made in the first part is weakened, in fact, that on accepting the second part, the first part is no longer acceptable, i.e., that the second part **corrects** the first part. The only reasonable explanation for this is, I think, that after all SP($\neg$A) and **must** A are semantically inconsistent if evaluated on the same information state. In other words: language users treat indirect information that A as good information, information that is not more compatible with the possibility of $\neg$A than the direct information A itself is: the seeming compatibility is in fact a case of pragmatic correction.

This argument, if plausible, has some strong consequences. Kratzer's theory is, as I indicated, in fact the only available option in a theory that builds information states out of possible worlds, more precisely, a theory that reduces truth on the basis of an information state to truth in the worlds in that information state, i.e., a theory that tries to define assertability conditions in terms of truth conditions. I

have argued that the distinction that underlies the difference between A and **must** A is not a distinction between good and less good evidence, but a distinction between good direct and good indirect evidence. As I have indicated here and have argued at length in Landman 1986, this is a distinction that you cannot make in a classical theory like possible world semantics; the fact that you can make the distinction in data semantics depends crucially on the fact that the partiality enters into the semantic recursion there. But then we have a strong argument in favour of the data semantic approach. The central notions of Stalnaker's and Kratzer's theories can be formulated as easily on information states as they can be formulated on sets of possible worlds. In fact, possible world semantics itself can be formulated in data semantics, if wanted, because total information states can be taken as possible worlds and accessibility relations can be defined on them (also, there is no principled reason to take $\leq$ to be the only accessibility relation on information states). But data semantics gives us more than possible world semantics, because it allows for a better account of the direct/indirect contrast.

Identity Statements And Epistemic Modals

Let me now turn to identify statements. Since identity is a two-place relation, the clause for atomic formulas that I have given earlier provides the basics of the semantics for identity statements:

$g \Vdash_s a = b$ iff $\langle g(a), g(b) \rangle \in i^+(=,s)$
$g \dashv\Vdash_s a = b$ iff $\langle g(a), g(b) \rangle \in i^-(=,s)$

g(a) and g(b) are objects in the model (in the domain of s). But what are objects? Are they classical, 'real' objects? What are real objects? Well, whatever a real object is, it is clear that if something is a real object, it is identical to itself and to nothing else. This means that if the objects in our models are real objects, the positive and negative extension of the identity are fixed in the following way:

$i^+(=,s) = \{\langle d,d \rangle : d \in D(s)\}$
$i^-(=,s) = \{\langle d,d' \rangle : d,d' \in D(s) \,\&\, d \neq d'\}$

The problem with this is that this makes = a **total predicate** on the domain of every information state. That is, for every information state s, it holds that if a and b are defined on s (that is if $g(a), g(b) \in D(s)$), then:

$g \Vdash_s a = b$ or $g \dashv\Vdash_s a = b$.

In other words, there is no way of **not** knowing whether a and b are the same or not.

Let us take the following drs, representing a Babylonian conversation concerning Hesperus and Phosporus: (this could be the drs of: *Something shines in the morning, something shines in the evening. Maybe they're the same, maybe not.*

D_4:

x y
x shines in the morning *y shines in the evening*
may $x = y$
may $x \neq y$

If an embedding g verifies the first two conditions relative to some information state s, then g(x), g(y) ϵ D(s). But that means that either $<g(x),g(y)> \epsilon i^+(=,s)$, or $<g(x),g(y)> \epsilon i^-(=,s)$. But if that is the case, then that is the case for every extension of s as well:

If $<g(x),g(y)> \epsilon i^+(=,s)$ then $\forall s' \geq s$: $<g(x),g(y)> \epsilon i^+(=,s)$
If $<g(x),g(y)> \epsilon i^-(=,s)$ then $\forall s' \geq s$: $<g(x),g(y)> \epsilon i^-(=,s)$

This just follows from the fact that our interpretations are growing (and given the interpretation of objects as classical objects, there is certainly not a way around that). But this means that if x and y are defined, **may** x = y is logically equivalent to x = y, and similarly **may** x $\neq$ y is logically equivalent to x $\neq$ y. Namely: clearly x = y implies **may** x = y. If x and y are defined on s then either g $\Vdash_s$ x = y or g $\Vdash_s$ x $\neq$ y. If g $\Vdash_s$ **may** x = y then for some extension s' of s: g $\Vdash_{s'}$ x = y. But then, given the above it is impossible that g $\Vdash_s$ x $\neq$ y, hence g $\Vdash_s$ x = y. The same argument for x $\neq$ y and **may** x $\neq$ y.

This shows that, if identity is a total predicate on every information state, the above drs D_4 is a contradiction. But clearly it is not, D_4 represents the information of someone who wonders about whether Hesperus and Phosphorus are one and the same object or not, and there is of course nothing incoherent in that kind of wonder.

There is only one way out of this problem: we have to make = a partial predicate. That is, we have to **weaken** the classical conditions:

$i^+(=,s) = \{<d,d>: d \epsilon D(s)\}$
$i^-(=,s) = \{<d,d'>: d,d' \epsilon D(s) \& d \neq d'\}$

Since I see no reason to remove the ordered pairs <d,d> from the positive extension of =, this means that:

$i^+(=,s) \supseteq \{<d,d>: d \epsilon D(s)\}$
$i^-(=,s) \subseteq \{<d,d'>: d,d' \epsilon D(s) \& d \neq d'\}$

with the normal condition that the positive and the negative extension of = on s do not overlap.

This means two things. In the first place, if d ϵ D(s) then $<d,d> \epsilon$ $i^+(=,s)$. So much is as before. But now there are objects d_1, $d_2 \epsilon$ D(s) of which we may find out that they are identical and of which we may find out that they are not identical, there are the objects in the gap of = on s. In other words, we cannot find out of d that it is not identical to d, but if d and d' are in the gap of = on s, then there may be an extension of s, s' such that $<d,d'> \epsilon i^+(=,s')$, and there may be an extension s'' of s where $<d,d'> \epsilon i^-(=,s'')$.

As I said before, of classical objects we know what they are identical to (themselves). So the things in our models are not classical objects.

(All this doesn't come as a surprise, of course, for readers familiar with intuitionistic logic. But note, for the interpretation that I will give in the next section, and its semantic consequences it is crucial that partial objects are the objects in D(s), and not their equivalence classes under identity on s, as in intuitionistic logic.)

We have to put some more restrictions on the identity predicate (for a more detailed account, see Landman 1986). First a definition:

> d and d' are **indiscernible on the basis of** s iff no property tells them apart on the basis of s (i.e., s does not put d in the positive extension of some predicate and d' in the negative extension of that predicate or vice versa). Else d and d' are **discernible on the basis of** s.

When are d and d' in the **negative** extension of = on s? Obviously, if s can tell d and d' apart:

> $<d,d'> \epsilon i^-(=,s)$ iff d and d' are **discernible** on the basis of s.

When are d and d' in the **positive** extension of = on s? Not if d and d' are indiscernible on the basis of s, not if s cannot tell them apart. Maybe some extension of s can. I take them to be identical on the basis of s if it is not possible to extend s to an information state that tells d and d' apart:

> $<d,d'> \epsilon i^+(=,s)$ iff for all $s' \geq s$: d and d' are indiscernible on the basis of s'.

Assuming, as I do, that every information state can be extended to a total one, that is, an information state where all predicates are totally defined (i.e., on every possible way of getting better informed, in the end you reach a state where all predicates are total), it now follows that = is an equivalence relation on every information state.

If the objects in our model are not real objects, what then are real

objects? Possible worlds, I have already said, correspond to total information states. Let w be a total information state (say, the information state corresponding to the real world). The real **objects** in w are the equivalence classes of the objects in D(w) under = on w.

In this way we reconstruct objects as their informational approximations (with respect to a certain standard of precision).

If one feels dissatisfied with this, one can of course live one's realist convictions by adding to the structures that I have described, domains of possible worlds and their individuals with a correspondence relation between the informational approximations and (some of) the real worlds and objects. (The situation is similar to that of event logic. The program of constructing moments of time out of events (see for instance Kamp 1979) does not forbid some (e.g., Hinrichs 1985) to add to the event structures a primitive domain of (real) moments of time).

Such an addition may satisfy some who feel that in the theory sketched here the solid ground of the real world is somehow disappearing under their feet. Whether there are compelling **semantic** reasons for such a modification remains to be seen, however.

Given this reinterpretation of the identity predicate, we can give a consistent interpretation to drs D_i:

D_i is true on the basis of s iff for some g:
$g(x) \in i^+$(shine in the morning,s) and
$g(y) \in i^+$(shine in the evening,s) and
for some $s' \geq s$: $\langle g(x),g(y)\rangle \in i^+(=,s')$ and
for some $s'' \geq s$: $\langle g(x),g(y)\rangle \in i^-(=,s'')$

and this is perfectly consistent.

We are, then, forced to the conclusion that objects at the level of discourse semantics, the level of information, are not classical objects, but partial objects, objects only partially determined by their identity conditions. How should we interpret such non-classical objects?

PEGS

What are pegs? Pegs are things to hang coats on and take them off again. They don't change themselves if you do that. The objects in our domains are pegs, informational pegs. Information states hang properties on pegs (and relations on tuples of pegs) in the process of information growth (and take them off again in correction) (in other words, information states ascribe properties to pegs). Since identity is

a relation, identity conditions are also hung upon pegs in this way (not just upon **a** peg, but upon pegs).

Pegs are informational objects, they are the trick that we use to overcome the problem of communication in a situation of partial information. We are commonly too badly informed to distinguish one real object from another that we know as much about (for instance, when someone reports to us: 'there's a man coming down the street. He wears a hat.'). **Pegs are objects we postulate in conversation as stand-in's for real objects. They are means of keeping track of what we talk about in information exchange.** As such they are postulated in shared information: we share the pegs we are talking about (because that is what they are for).

We cannot reach real objects. So we postulate pegs as stand-ins for real objects. When one person postulates a peg, other people can continue, and ascribe properties to the same peg. We talk about them as if they have independent existence, existence outside of us, like real objects.

It should be clear that I am taking a strongly pragmatic view on the content of the model-theoretic information structures that I am using. Information structures are pragmatically postulated, intersubjective structures "in between" mental representations and the world. The key to their existence is agreement between language users. Neither a strongly mentalistic, nor a strongly realistic view on those structures is particularly suited. Pegs are not mental objects, but agreed upon **interpretations** of those, and their nature as postulated objects does not make them very real. This is not to say that the theory can do without either mental representations or the world. Though it is a pragmatic choice **which** pegs are actually postulated in a conversation, and pegs are conventional in that way, this does not mean that the **structure** of the level of information is. In other words, though information structures form a medium in which we can make relatively arbitrary choices, this does not mean that the medium is arbitrary in any way, in fact we would expect the medium to be strongly determined both by our cognitive apparatus and the world.

But pegs are pragmatic objects. It is the context in a conversation, and its standard of precision that determines what distinctions are made, i.e., that determines the domain of pegs, and that manipulates domains of pegs: context change can make things that were discernable before indiscernible by adding new pegs (for examples, see Landman 1986).

Pegs are postulated objects, whose function is to keep track of objects in information states. When, in an introduction to semantics class, we talk about Hesperus and Phosphorus, we talk **as if there are**

two objects. We use the plural **they**, even when we say that **they** are identical, and when we say that they are not, on the basis of the Babylonian information. It is, I think, important to realize this. For instance, when Kripke (Kripke 1979) presents his puzzle about belief, he argues that the puzzle does not rely on substitution, but only on the uncontroversial assumption that in representing a French sentence in English we can translate 'Londres' by 'London.' This assumption is uncontroversial indeed in normal circumstances, but precisely **not** in the case that Kripke discusses. In order to translate a French name into its English equivalent it does not suffice to assume that they stand for the same object, but we need the further assumption that the original and its translation stand for the same peg, the same conversational object. And this is not the case in Kripke's case; indeed, the case is constructed in such a way that precisely that condition is violated: 'Londres' and 'London' stand for two pegs, referring to the same real object. When Frege talks about 'Venus,' 'der Morgenstern,' 'der Abendstern,' he talks about them as if they are three things, the text postulates three pegs. And, of course, a translator cannot say: 'Oh, it's so confusing, those different expressions, let me all translate them with Venus.' The situation is different if an astronomy book uses those three names as stylistic variants. In that case there is only one peg, and no problem arises if a translator decides to eliminate that variation. (For more discussion of Kripke's puzzle, see Landman 1986.)

Still, when we talk about the world-view of the Babylonians and introduce two pegs for Hesperus and Phosphorus, we use **the same** pegs when we describe their information, as we do when we compare it with ours. This is crucial for correction. Only in this way can we hope to convince the Babylonians that they have wrong information, in other words, that they have attributed the wrong properties to these same pegs. In this sense **pegs are objective** (or conversationally uniform). In a picture as seen on next page.

We see indeed that pegs perform their primary function: we keep track of Hesperus and Phosphorus through different information states.

The theory as developed here is meant to combine three seemingly incompatible theories:

(1) Pegs are like (Fregean) senses, **conceptual** objects. Though they are postulated by us, they are postulated as independent of us. They are more intensional, conceptual, though, than intensions in possible world semantics. If $\Box(h = p)$, then h and p have the same intension. However, as I have indicated above, even if $\Box(h = p)$ is true, we can assume that h and p are interpreted as different pegs. This is so, precisely because pegs are conversationally postulated objects.

(2) For the same reason, pegs are like (Husserlian) intentional objects, objects of postulation, the objects as we talk about them, as we intend them, the objects we aim at.

It will be clear that given this, pegs have a lot in common both with Castaneda's guises and with Kaplan's vivid names, to mention a few. There is a third aspect, however, that I'd like to stress:

(3) Pegs are like (Kripkean) rigid designators. Consider the following quotation from Kripke 1972: "A possible world isn't a distant country that we are coming across, or viewing through a telescope. . . . A possible world is **given by the descriptive conditions we associate with it**. . . . Why can't it be part of the **description** of a possible world that it contains **Nixon** and that in that world **Nixon** did not win the elections? . . . There is no reason why we cannot **stipulate** that, in talking about what would have happened to Nixon in a certain counterfactual situation, we are talking about what would have happened to **him**" (Kripke 1972:267).

The view on counterfactuals that seems to fit this passage very well is that we take Nixon in our world and bring **him** to different worlds, by changing the facts. The fact that the name 'Nixon' denotes in those worlds the same person as it does in ours, then, is a direct consequence of this interpretation of the counterfactual: we keep track of Nixon in the process that brings us from our world to counterfactual ones.

The analyses of counterfactuals that are most faithful to this idea are those of Veltman 1976 and Kratzer 1977. The basic idea is that, for a counterfactual 'If Nixon had lost in 1968, then . . .,' you first go back in information till "Nixon lost in 1968" becomes compatible, and then

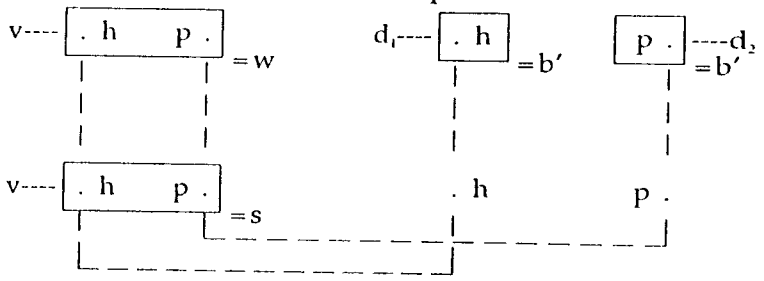

our information s

The Babylonian information b (as described by us)

you evaluate the conditional as an indicative conditional (i.e., you check whether if you add the antecedent to those information states, that commits you to the consequent). Pegs, as I have stressed, are primarily means of keeping track of objects through information change. In my perspective, if, in evaluating the counterfactual, you go to other information states, by removing properties that our present information state ascribes to the peg Nixon, it is still **that peg** that you take to different information states. In this sense, you keep track of Nixon in the course of information change, and in this sense *Nixon* is a rigid designator.

One might say that I am perverting the notion of rigid designation. One might say the following: 'Okay, you're keeping track of Nixon, by taking that peg to different information states. But your theory allows us to take the peg n(ixon) to an information state on the basis of which n = j(ane fonda) is true. You can't possibly do that on Kripke's view. So there is clearly a sense in which your notion of rigid designation differs from Kripke's. Or to say it differently, Kripke's notion of rigid designator has philosophical bite, while yours does not."

I tend to agree with this criticism, but I also tend to prefer my interpretation. The judgments on the truth and falsity of counterfactuals in as far as they bear on rigid designators are a consequence of two aspects of the semantics (and/or pragmatics) of counterfactuals. The first is **minimality**. The assumption shared by most theories of counterfactuals is that in order to evaluate a counterfactual you go to information states (worlds) that are **minimally** different from the present one in that the negation of the antecedent doesn't hold. Although technically identity is a relation like others, this does not mean that language users are not aware that its semantic nature differs from a relation like, say, **love**. It is not at all implausible to assume that in general you stay closer to this world if you change the extension of relations like **love** than if you start tinkering with the ontology. It seems to be a plausible constraint on the minimality involved in counterfactual change that counterfactuals will not remove identity conditions **unless** they are explicitly forced to (I will come back to this).

It is this minimality that makes us expect that in discussing whether Nixon could have been a movie star, we will try to change Nixon's career decisions at various times of his life, to see whether we end up in Hollywood, rather than trying to make him identical to Jane Fonda.

Still, this doesn't answer the question: Could Nixon possibly have been Jane Fonda? This is the real question at issue: Kripke's answer, the answer that we have learned to accept, is a strong no. This is

where the second aspect of the semantics of counterfactuals comes in. This second aspect is the context dependency we have seen before. Kripke is very clear about what he is concerned with: he is asking a metaphysical question: given the way the world is and the fact that Nixon is not Jane Fonda, could that fact have been different? Given that the answer to that metaphysical question is no, it is clear that if names are rigid designators the counterfactual process that leads us through the sphere of metaphysically possible worlds will never identify the two.

The modal base, then, with respect to which these counterfactuals are evaluated is one that contains our metaphysical assumptions. I have no problem with the assumption that for a "metaphysical modal base" (that is, if the information state and the accessibility relation tell us what is possible in view of our metaphysical assumptions) there is a constraint that the identity relations are the same in all accessible information states (or at least in all total ones). But a metaphysical interpretation is not the only one that counterfactuals can get. It is at least as natural (and quite a bit more common, I would say) to give counterfactuals an epistemic interpretation, relating them to an epistemic modal base and, as Kripke has taught us, we have to account for the fact that statements that may be metaphysically necessary can be 'epistemically contingent.' The whole point of the information semantics for identity statements that I have given is to account for this epistemic contingency, to account for the fact that statements like *maybe Hesperus is Phosphorus, maybe not* and questions like *Is Hesperus Phosphorus?* are **contentful**, meaningful, and that it makes sense to raise them, and similarly that one can be surprised about the answer. Above, I formulated a minimality condition on counterfactuals that I repeat here: counterfactuals will not remove identity conditions **unless** they are explicitly forced to. My claim is that for counterfactuals, we cannot understand their meaningfulness if we do not allow for this 'unless we are explicitly forced to,' if they are related to an epistemic modal base.

Consider the following sentence:

(38) If Hesperus had not been Phosphorus, then Babylonian science would be more respected nowadays.

If you would say this sentence at a party to me, I would disagree with you, because I think that Babylonian science is quite respected nowadays, and the Hesperus/Phosphorus business would not have made a difference.

In disagreeing with you, I take your sentence as contentful (as contentful as (39):

(39) Suppose we would find out tomorrow that Hesperus is not Phosphorus after all. Would Babylonian science be more respected?

We can perform such thought experiments and quarrel about their outcome. All that would be impossible in possible world semantics. (38) is vacuously true, because $h \neq p$ is neccessarily false, false in all possible worlds.

If we assume that Hesperus and Phosphorus are pegs that are identical on the basis of our present information, then (38) and (39) instruct us to go back in information to states where their identity conditions are no longer certain and see what happens if we add the information that they are not identical. Although we may not be able to tinker with the ontological furniture of the world, we are certainly able to imagine that we wouldn't have certain information that in fact we do have, and we are able to imagine what we would do if our information were different. It is this capacity that makes counterfactuals like (38) meaningful, and our semantic theory has to be rich enough to account for this.

What we need, then, is a more intentional theory of meaning (both in the philosophical sense and in the technical sense of Thomason 1980), one that does not lose aboutness, but has a place for epistemic contingency. This will necessarily be a theory that has a less direct relation between meanings and extensions than the notion of an intension, as a function from possible worlds to extensions, gives us. In other words, this has to be a theory that in fact does not equate meaning and necessity.

It seems to me that the information perspective that I have sketched can form an interesting guide line for such a theory and that the theory of pegs that I have outlined here can form a natural part of it.

ACKNOWLEDGMENTS

*I would like to thank Nirit Kadmon for stimulating discussions and comments that made me write the present paper and write it in its present form. Further, discussions at various times with Gennaro Chierchia, Angelika Kratzer, Barbara Partee, Craige Roberts, and Ray Turner were very helpful. This paper has further benefitted from the comments of one of the anonymous referees.

Parts of the present paper were presented in various talks that I gave at the University of Massachusetts at Amherst and at Cornell University. I would like to thank the audiences of those talks for their comments.

Part of the research for this paper was done while I held a research position at the University of Massachusetts, financially supported by Barbara Partee's grant "Foundations of Semantics" from the System Development Foundation (SDF). This support is gratefully acknowledged.

REFERENCES

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*, Diss., GLSA, Amherst: University of Massachusetts

Hinrichs, E. (1985). *A Compositional Semantics for Aktionsarten and NP Reference in English*, diss., Ohio State University

Kamp, H. (1975). Two theories about adjectives. In Keenan (ed.), *Formal Semantics for Natural Language*. Cambridge, Eng.: Cambridge University Press

— (1979). Events, instants, and temporal reference. In Bauerle, Egli, and von Stechow (eds.), *Semantics from Different Points of View*. Berlin: Springer

— (1981). A theory of truth and semantic representation. In Groenendijk, Janssen, and Stokhof (eds.), *Formal Methods in the Study of Language*. Amsterdam: Mathematical Centre Tracts. Reprinted in Groenendijk, Janssen, and Stokhof, (eds.) (1983), *Truth, Interpretation, Information*, GRASS 2. Dordrecht: Foris

Kaplan, D. (1979). On the logic of demonstratives. *Journal of Philosophical Logic* 8

Karttunen, L. (1972). Possible and must. In Kimball (ed.), *Syntax and Semantics 1*. New York: Academic Press

Kratzer, A. (1977). What 'must' and 'can' can and must mean. *Linguistics and Philosophy* 1

— (1981). The notional category of modality. In Eikmeyer and Rieser (eds.), *Words, Worlds and Contexts*. Berlin: de Gruyter

— (1986). On conditionals. In Farley, Farley, and McCullough (eds.), *Papers from the Parasession on Pragmatics and Grammatical Theory*, CLS 22

Kripke, S. (1972). Naming and necessity. In Davidson and Harman (eds.), *Semantics of Natural Language*, Dordrecht: Reidel

— (1979). A puzzle about belief. In Margalit (ed.), *Meaning and Use*. Dordrecht: Reidel

Landman, F. (1986). *Towards a Theory of Information*, GRASS 6. Dordrecht: Foris

Landman, F. (1987). A handful of versions of discourse representation theory, ms, Cornell University

Lewis, D. (1975). Adverbs of quantification. In Keenan (ed.), *Formal Semantics for Natural Language*. Cambridge, Eng.: Cambridge University Press

— (1979). Scorekeeping in a language game. In Bauerle, Egli, and von Stechow (eds.), *Semantics from Different Points of View*. Berlin: Springer

Partee, B. (1987). Quantifying over contexts, paper presented at the Sixth Amsterdam Colloquium, April

Roberts, C. (1987). *Modal Subordination, Anaphora, and Distributivity*, diss., GLSA, Amherst: University of Massachusetts

Stalnaker, R. (1968). A theory of conditionals. In Rescher (ed.), *Studies in Logical Theory*. Oxford: Blackwell

— (1974). Pragmatic presuppositions. In Munitz and Unger (eds.), *Semantics and Philosophy*. New York: New York University Press

— (1979). Assertion. In Cole (ed.), *Syntax and Semantics* 9. New York: Academic Press

— (1984). *Inquiry*. Cambridge, MA: MIT press

Thomason, R. (1980). A model theory for propositional attitudes. *Linguistics and Philosophy* 4

Veltman, F. (1976). Prejudices, presuppositions and the theory of Conditionals. In Groenendijk and Stokhof (eds.), *Proceedings of the First Amsterdam Colloquium*, Amsterdam Papers in Formal Grammar 1. University of Amsterdam: Dept. of Philosophy

— (1981). Data semantics. In Groenendijk, Janssen, and Stokhof (eds.), *Formal Methods in the Study of Language*. Amsterdam: Mathematical Centre Tracts. Reprinted in Groenendijk, Janssen, and Stokhof (eds.), (1983). *Truth, Interpretation, Information*, GRASS 2. Dordrecht: Foris

— (1985). *Two Logics for Conditionals*, diss., University of Amsterdam

— (1986). Data semantics and the pragmatics of indicative conditional sentences. In Traugott et al. (eds.), *On Conditionals*. Cambridge, Eng.: University Press