PEGS AND ALECS

Fred Landman
Department of Philosophy
University of Amsterdam

ABSTRACT

A major problem in semantics is the question how identity statements can be informative. To answer this question we have to determine what the status is of the objects that language users talk about when they exchange information. It is argued that the assumption of discourse representation theory, that these objects are variables in some representation, leads to problems. To cope with these problems, a theory of pegs is developed, partial objects at an intermediate level of information, to which the partial information states of language users ascribe properties, and that language users can keep track of in the process of information growth. This theory is applied to notorious problems of identity, like the morningstar-paradox, Kripke's puzzle about belief, and the paradox of the hooded man. Within the theory of pegs, an analysis of donkey-sentences is given, that resembles the analysis of discourse representation theory, except that it is not based on variables. To this end, alecs are introduced, pegs which, relative to some information state, can play the role of all pegs with certain properties.

1. Partial objects and identity

This paper is concerned with the analysis of identity statements in frameworks of partial information, and with the status of partial objects. More in particular, it contains a comparison between discourse representation theory and data semantics on these issues, and it develops a theory of partial objects within the latter framework.

The main problem it addresses is the proper treatment of epistemic puzzles of identity, like the well-known Hesperus-Phosphorus puzzle. The problem can be formulated in general as follows: what and where are the objects we talk about? The Babylonians talk about Hesperus and Phosphorus, and agree that they talk about two things, i.e. they talk as if there are two. Yet, if the objects they talk about are private representations of real objects in their minds, then there are four objects (two for each speaker). On the other hand, if they are real objects in the world, there is only one object. What is the level of objects talked about, where there are two? Further, our information about the objects we talk about may grow, but we assume that they are the same objects, before and after our conversation: we can keep track of the objects we talk about in the process of information growth. Finally, the Babylonians are not aware of the identity of Hesperus and Phosphorus, and they regard the information that this is the case as completely new to them, as epistemically contingent.

It is notoriously hard for theories to capture all these things at the same time.

Discourse representation theory is a collective name for two formally equivalent theories, that of Kamp [1981] and that of Heim [1982].

On Kamp's view, language users make representations of the utterances they hear and add these to representations of earlier discourse. These representations contain discourse markers, marking the objects talked about; predicates, marking the properties ascribed to them; and logical structure. Representations get a semantics because language users, so to say, compare these representations with the world: embedding functions map discourse markers onto objects in the world, satisfying recursive embedding conditions.

Heim's theory distinguishes a level of discourse syntax, containing discourse markers (which are variables), and a level of discourse semantics, containing discourse referents ('file cards'). The embedding conditions are here part of the discourse semantics: the interpretation of a (discourse-) syntactic representation is a file, a set of discourse referents and a set of embedding conditions.

The theories are equivalent, because, for Heim, discourse markers are *variables*, and discourse referents are (classical) *variable meaning*. As in classical logic, a variable meaning (be it the index of the variable, or that function that assigns to every assignment function the value of that variable for it) can, without problems, be identified with the variable itself, and from that the equivalence follows.

So, in discourse representation theory, the objects talked about are variables, linked through embedding conditions to real objects in the world (the model). There are two ways that identity conditions can be dealt with: either—as one would expect—through the embedding conditions, or directly at the level of representation. However, there are problems with both.

As for the first strategy, though the Babylonians are not aware of the identity, they do point at an object in the sky. The embedding conditions of their discourse should map both Hesperus and Phosporus on what they point at. But then, adding an identity condition to the discourse will not reduce the set of embedding conditions, and will provide no new information. (Note that it does not help to replace the world in which their discourse is embedded by a set of possible worlds or a situation (a chunk of the world), since what they point at twice will be in that chunk of the world every time they point at it, and it will be constant through that set of possible worlds.)

The second strategy runs into different problems. Variables are syntactic expressions, symbols. In that, they are like words. But symbols have trivial identity conditions: a symbol is identical to itself and to nothing else (that is, *as* a symbol). You cannot make two symbols into one, and hence you cannot identify Hesperus and Phosphorus.

I will argue that, in order to make the identity conditions epistemically contingent, we cannot regard the objects we talk about as real objects in the world, nor can we take them to be symbols in some representation; they are semantic objects with non-trivial identity conditions, at an intermediate level of information, that language users postulate and share: *partial objects* that can merge in the process of information growth.

In agreement with Heim, I will defend the view that the objects we talk about are ingredients of a level of discourse semantics, and that, as such, they should be understood as 'real objects in disguise'. But they are not variables or variable meanings, but partial objects.

I will very shortly discuss two approaches to partial objects.

Parsons [1981] takes them to be sets of properties. Hesperus is a set of properties that Hesperus has, Phosphorus a set of properties that Phosphorus has; the two need not be the same. However, Parsons combines this with a classical analysis of identity. This has the consequence that, if the information grows, we cannot say that we get better informed about Hesperus and Phosphorus. We can only say that the old objects Hesperus and Phosphorus are replaced by new ones. We cannot keep track of Hesperus and Phosphorus in the process of information growth and in the end, the discovery that Hesperus is Phosphorus is turned into the linguistic decision to choose new referents for the names Hesperus and Phosphorus, like choosing a penname.

Scott [1980]'s domain theory is somewhat similar to Parsons' theory: partial objects are constructed as proper filters of properties, or propositions, in an information system. An *information system* is a structure of propositions (also called possible facts) based on two informational notions: that of information containment and that of incompatibility of information. Both are relations between propositions. Formally, we take an information system to be a meet-semilattice, a set partially ordered by information containment and closed under combination of information, with a minimal element: the impossible proposition, containing incompatible information. Two propositions are incompatible if their combination is the impossible proposition.

The advantage over Parsons' theory is that the information perspective allows a non-classical, monotonic analysis of identity, like: two objects (proper filters) are identical if they can only grow towards one and the same total object (ultrafilter). Such a non-classical analysis of identity is essential for a theory that takes partial objects seriously: it is classical identity that makes an object identical to itself and to nothing else, that makes an object total. On Parsons's approach, information growth comes down to replacing total objects by larger total objects: hence, Hesperus on one stage of the information is a different object from Hesperus at some other stage, and if Hesperus and Phosphorus are different sets of properties, they are non-identical. The present, non-classical, analysis makes the identity of Hesperus and Phosphorus contingent: they can be identical even if they are different partial objects (if only they grow towards the same total object). In this respect, we have gained some advantage. But not enough. We still cannot keep track of Hesperus in the process of information growth. As long as Hesperus can grow towards different total objects, it is a different object at different stages of the information. And this has the unfortunate consequence that it is not only the identity of Hesperus and Phosphorus that is contingent, but that of Hesperus and Hesperus itself as well: as long as Hesperus can grow towards different total objects, the statement that Hesperus is Hesperus is undefined.

2. Data semantics

In the next section I will discuss the analysis of these problems in data semantics. Here I will very shortly introduce the basic principles of that theory.

On the ontological side, we assume a set of 'objects' D, an information system of propositions $\mathcal{I}$, a set of properties $\mathcal{P}$, and an interpretation function ι, interpreting constants directly as objects, and predicates as properties. Properties are certain functions mapping n-tuples of objects onto propositions.

Information states are defined as sets (proper filters) of propositions. This gives us notions like 'extension of inmformation', 'total information' (relative to the model), etc.

The main partiality comes in with the semantics: instead of a recursive definition of truth conditions, the semantics gives a recursive specification of conditions for truth ($\Vdash$) and falsity ($\dashv\Vdash$) *on the basis of* an information state **s** (verification/falsification conditions, assertability conditions, or conditions of evidence). Since properties map objects onto propositions, atomic formulas can be interpreted as propositions. The semantic behaviour of modals and conditionals is defined in terms of possible extensions of information states. Examples of clauses are:

$\mathbf{s} \Vdash P(c)$ iff $\iota(P)(\iota(c)) \in \mathbf{s}$
$\mathbf{s} \dashv\Vdash P(c)$ iff $\iota(P)(\iota(c))$ is incompatible with some proposition in $\mathbf{s}$

$\mathbf{s} \Vdash \mathbf{may}\varphi$ iff for some $\mathbf{s}' \supseteq \mathbf{s}$ $\mathbf{s}' \Vdash \varphi$
$\mathbf{s} \Vdash \mathbf{must}\varphi$ iff for all $\mathbf{s}' \supseteq \mathbf{s}$ there is an $\mathbf{s}'' \supseteq \mathbf{s}$ such that $\mathbf{s}'' \Vdash \varphi$

For reasons given in Landman [1986], these clauses for modals are too simple, but they suffice for our purposes here. Unfortunately, this is not so for the conditional: later in this paper I will use the rather involved analysis of conditionals that was developed in Landman [1986]. However, I cannot here give all the definitions that it uses, so I will describe the semantics of the conditional informally here, and refer to Landman [1986] for the formalization:

$\mathbf{s} \Vdash (\varphi \rightarrow \psi)$ iff for every way of extending $\mathbf{s}$, if the first thing that happens to the antecedent φ there is that it becomes true, then right after that has happened, the consequent ψ must be true

A crucial consequence of adopting assertability instead of truth as the basis of the recursion is that modals (like **may**) can be instable under extension of information (true on the basis of limited evidence, false later on).

The main concern of the present paper is: what are these 'objects'? They are to be partial objects, but in what respects? If we don't assume them to be real objects in the world, nor private objects in our minds, then what are they? It is time to introduce the heroes of this paper: they are pegs.

3. Pegs

What are pegs?
Pegs are things to hang your coat on. Or your hat. They change their appearance, if you do that. But we find no reason in that to say that it is not the same peg. Of course, real pegs are physical objects, and we hang hardly anything on them but cloths. Our pegs are informational objects, and the things we hang on them are properties.

As always there has to be someone who hangs the cloths on the peg, or takes them off again. Pegs are not made with the hats already on them, nor is it impossible to take the hats off them, or else they would be worseless pegs, that's simply not what they are for. It is essential for being a peg that you can hang coats on them, and someone has to do it. In our case, it is the information available to us that hangs the properties on the pegs: our cloak-room attendants are the information states.

Properties map pegs on facts; information states contain facts. Suppose we have a peg p and a property P. If our information state does not yet contain the fact $P(p)$ then it has not yet hung the property on the peg. It has, if it grows into a state that does contain that fact. They are not very special pegs, actually, except that you can hang a lot on them. Our cloak-room attendent is of the old-fashioned kind, however. It abhors to hang two coats on one peg of which the colours clash, and so it refuses to do that. We will see that it has other peculiar character traits, in due course.

Now what about the following questions:
When does a peg have a certain property?
Do pegs exist?
How many pegs are there?

It is important to realize that—even if we want to make ourselves feel liberal—these questions are *completely irrelevant*. I tend to be less liberal in this matter, and think that they are not just irrelevant, but *nonsensical*. It is a nonsensical question to ask whether a peg has a certain property: pegs don't *have* properties; properties are *ascribed* to them by information states. But don't we say that an object in the world has a certain property and couldn't that object be a peg? To say that an object in the world is a peg which has a certain property should be interpreted as to say that the complete information state which we assume to be the best possible approximation of the real world *ascribes* that property to that object. It is a nonsensical question to ask, out of the blue, whether a peg exists; information states ascribe existence to pegs. Similarly, it is nonsensical to ask for the cardinality of the set of pegs, *because pegs do not have identity conditions*, identity conditions are no better than properties, they are ascribed to pegs

by information states; pegs can be identified by information states. If we want to ask how many objects there are according to a certain information state, then we should count the number of distinct sets of pegs that are identified by that information state. But the cardinality of that domain changes in the course of information growth. But of course, interpreted in this way, it is a legitimate question to ask, say, how many objects there are in the world, because—as we will see—on total information all identity conditions are fixed, and we ask for the number of *distinctions* that this information state makes, that is, the number of equivalence classes of pegs under identity.

To formulate this matter somewhat differently: I refuse and reject it as nonsensical to regard pegs as classical, static objects, that have properties, that do or do not exist, and that are identical with themselves only; *by squeezing them into the armour of classical logic in this way, we try to understand a dynamic process as a static one.* Compare it with a movie. A movie is a dynamic entity which gains its sense by the things happening in it. We can stop the movie at a certain point, freeze the action at a certain scene; that can even help us to understand the movie better. But taking a classical view on the movie, is taking an *external* view on it, regard it as an ordered set of distinct images. Doing that might serve some purposes, but we cannot understand the *meaning* of the movie in that way, because its meaning lies in the action, in the things that are happening in it. It is impossible to understand a movie if, while watching it, you seriously regard it as a long sequence of distinct pictures. It is only by ignoring that aspect of the movie, by stepping in it, that is, by taking an *internal* view, that we can make sense of it.

Information exchange is a dynamic process. In modelling it, we make a classical reconstruction of it, we use sets, classical model-structures, as the marble out of which we sculpt our reconstruction of this process, the statue which depicts the process. We can try to understand the depicted process completely in terms of the pieces of marble used, the material used to model the process, but we shouldn't; it may give us a better understanding of our reconstruction of the process, but not of the process itself.

Information states hang properties on pegs. They refuse to hang incompatible properties on one peg. Being partial objects themselves, they can be very short-sighted, in fact, they can only distinguish one peg from another by the properties that hang on them: they can tell them apart if one of them has a property incompatible with some property that the other one has. So there can be lots of pegs that an information state cannot tell apart. One might think that the customers will complain about that, and say that they don't get the right hat, because it comes from a different peg, but properties are very special hats, even if they come from different pegs, they are the same hat, so most customers don't complain. In fact, the only trouble makers are the classical logicians: they are the ones who insist to get their hat from the same peg, and in that they rudely force our information state to strain its eyes more than is justified.

For us, on the other hand, precisely the fact that sometimes pegs cannot be told apart on the basis of an information state forms the *key* to their nature: *the essence of partial information is that it cannot justify certain distinctions, and the decision about the identity of certain pegs is a prime example of that.*

Let us give some definitions:

d_1 is **discernable from** d_2 **on the basis of** information state **s** iff there is a property P such that:
either $P(d_1) \in \mathbf{s}$ while $P(d_2)$ is incompatible with some fact $f \in \mathbf{s}$
or $P(d_2) \in \mathbf{s}$ while $P(d_1)$ is incompatible with some fact $f \in \mathbf{s}$

Actually, we should generalize this to n-place properties where d_1 and d_2 are among the arguments, but this is straightforward, and we won't bother.

d_1 and d_2 are **indiscernable on the basis of s** iff they are not discernable on the basis of **s**

So d_1 and d_2 are indiscernable on the basis of the information if they don't have incompatible properties. We can also give a stronger notion of indiscernability: d_1 and d_2 are strongly indiscernable on the basis of the information if they have the same properties on that basis. We can distinguish various degrees of strength here. We can let d_1 and d_2 be indiscernable only with respect to the properties involved in the information, or let them be indiscernable also with respect to those properties that they *must* and *may* have on the basis of the information. It is the latter notion that we will define and use here:

d_1 and d_2 are **strongly indiscernable on the basis of s** iff for all properties P : $\mathbf{s} \Vdash P(d_1)$ iff $\mathbf{s} \Vdash P(d_2)$ and $\mathbf{s} \dashv\Vdash P(d_1)$ iff $\mathbf{s} \dashv\Vdash P(d_2)$ and
$\mathbf{s} \Vdash \mathbf{must}\, P(d_1)$ iff $\mathbf{s} \Vdash \mathbf{must}\, P(d_2)$ and $\mathbf{s} \Vdash \mathbf{may}\, P(d_1)$ iff $\mathbf{s} \Vdash \mathbf{may}\, P(d_2)$
Note that the fact that on the basis of our information they *can* have property P, does not necessarily mean that they will have P in the *same* extensions of the information, that would again be a stronger notion of indiscernability. Here we will be satisfied with calling d_1 and d_2 strongly indiscernable if our information ascribes the same properties to them; further if the properties that they must have on the basis of our information are the same, and similarly for the properties that they may have.

For later use we will also introduce a *one way* version of the latter notion:

d_1 is an **indiscernable approximation of** d_2 in s iff for all properties P : if $\mathbf{s} \Vdash P(d_1)$ then $\mathbf{s} \Vdash P(d_2)$ and if $\mathbf{s} \dashv\Vdash P(d_1)$ then $\mathbf{s} \dashv\Vdash P(d_2)$ and
if $\mathbf{s} \Vdash \mathbf{must}\, P(d_1)$ then $\mathbf{s} \Vdash \mathbf{must}\, P(d_2)$ and if $\mathbf{s} \Vdash \mathbf{may}\, P(d_1)$ then
$\mathbf{s} \Vdash \mathbf{may}\, P(d_2)$

Now we can see why indiscernability contains the key to the proper understanding of pegs. Suppose that you don't have information at all, that $\mathbf{s} = \emptyset$. *Then all pegs are indiscernable.* Pegs get discernable only through information growth. So pegs can also be called *indiscernable objects*, because in the beginning, there were indiscernables.

So for some pegs your information may contain enough facts to tell them apart, but there can be lots of pegs, which are indistinguishable as far as your information is concerned. If we think about what it means for an identity statement to be true or false on the basis of an information state, we see that the *falsification* condition is straightforward:

$\mathbf{s} \dashv\Vdash (c_1 = c_2)$ iff $\iota(c_1)$ and $\iota(c_2)$ are *discernable* on the basis of s

We will, of course, *not* define $(c_1 = c_2)$ to be true on the basis of s if their interpretations are indiscernable: though s does not contain enough information to distinguish between those two, it can very well be that s admits the possibility that it grows into an information state that *does* contain enough information to distinguish between them. However, as soon as *that* possibility is eliminated, we *can* say that we have enough information to regard them as identical:

$\mathbf{s} \Vdash (c_1 = c_2)$ iff for no $\mathbf{s}' \supseteq \mathbf{s}$ it holds that $\iota(c_1)$ and $\iota(c_2)$ are discernable on the basis of s'
We will also call two *pegs* identical on the basis of s if they are indiscernable on the basis of every extension of s.

Before we go on, let us briefly return to the question: 'how many objects are there in the world?' and a related question: 'to which real object in the world does a peg refer?'.

Let w be a total information state, a maximal approximation of a possible world. The following holds for w:
If d_1 and d_2 are indiscernable on the basis of w then they are identical on the basis of w.
If a total information state does not contain a property that can distinguish between d_1 and d_2 then of course there is also no extension of it which does.
We call **the referent of** d **in** w the equivalence class:

$$\{d' : d' \text{ and } d \text{ are identical on the basis of } \mathbf{w}\}$$

The **real objects** in the world w are the equivalence classes under identity on the basis of w. So d_1 and d_2 refer to the same object in the world if they have the same referent in it. Also, the number of real objects in the world is the cardinality of the set of objects in the world. So we can have objects in our theory who behave like 'the real guys': they are constructions out of pegs under the assumption that our information is total.

We will put this notion of object to use later.

Note that we will often be in an information state s such that $\mathbf{s} \Vdash \mathbf{may}\,(a = b)$ and $\mathbf{s} \Vdash \mathbf{may}\,(a \neq b)$. Of course, once s has grown into an information state s' where $\mathbf{s}' \Vdash (a = b)$ it also holds that $\mathbf{s}' \dashv\Vdash (a \neq b)$

This, basically, is what pegs are. They are the objects we assume in conversation, and which we follow through information growth. On very partial information they are indiscernable from other objects to which they may be identical, but also from which they may be discernable in the end. Even though we may ascribe wrong properties to them, it is an *objective* matter to which object in the real world they refer and also what properties they have there: if our information state ascribes to something we point at a property which, after some consideration, we agree it doesn't have, we don't substitute a different object for it, we don't pretend that we were pointing at something different, but we change our information state to a state which ascribes different properties to it; we *correct* our information, that is, we transcend from an information state that was not part of the total information state w_0 which approximates the real world, to one that is.

I assume with Heim that there is a level of discourse semantics in between the level of discourse syntax and the real world, in fact that intersubjective level is our main level of interpretation, the real world plays hardly a semantic role at all, it is there as a regulative idea which sets our standard of correctness. In conversation we build up a structure or sequence of logical forms. These contain markers and predicates which are interpreted as pegs, the objects we talk about, and properties. Two different markers, distinct symbols, can be interpreted as the same peg, if our conversation assumes that we talk about only one object (as I argue in Landman [1986], there can be syntactic reasons for introducing distinct markers in the discourse, even if there is only one object talked about). Logical forms are evaluated with respect to an information state, the direct context, containing the information we agree upon. This information state is a set of facts, pegs are not in it, but each information state uniquely determines the domain of those pegs about which it contains information. Accepting new information leads to extension of the structure of logical forms and update of the information state, a transition from one information state to another. The new information state can contain more facts about pegs which were already in the domain of the old information state, as well as contain facts about a larger domain of pegs, introduce new objects to be talked about. The main difference between Heim's theory and mine lies in the relation between discourse syntax and discourse semantics (an embedding versus just a homomorphism) and in the nature of discourse referents (variables versus indiscernables). We will see later in this paper that this invites a difference in the analysis of complex discourses as well.

4. Strangers and acquaintances

We will now discuss some wellknown semantic puzzles.

Let us first return to Babylonia. We said that three things were very clear: the Babylonians talked about the same thing, when they pointed at Hesperus; the thing they talked about is the same thing before and after their conversation; and they are not aware that Hesperus is the same heavenly body as Phosporus. It is clear that if we assume Hesperus and Phosphorus to be pegs, that get their disguises through the information state that the Babylonians share, we can account for these three things. Though there may be four objects in their heads altogether and though there is only one object in the world, their information state ascribes properties to two pegs. And even if they have different information, even if there is not an information state (or part of it) which they share, their information states ascribe properties to those two pegs, in the situation described it ascribes the property of being visible at the evening sky to Hesperus and the property of being visible at the morning sky to Phosporus. Similarly, through information growth during the conversation, more properties are ascribed to both these pegs, but they stay the same peg: pegs are not locked in information states in the sense that every new information state is a new slide with new objects in it; we can follow pegs through information growth (and information change).

It is precisely this that distinguishes the present theory from those that identify partial objects with sets of properties: sets of properties *are* locked in information states: either you follow a particular set of properties in larger information, and you regard that particular set as an object distinguishable from its supersets, or you follow larger and larger sets of properties, but then you have no way along those paths to choose one way rather than another, i.e. to identify it with the one possible extension rather than another. Pegs don't have that problem because they choose their own way. If there are two possible sets of properties which a peg may adopt, it will choose one information state to adopt the one set, and

another to adopt the other; *within* an information state a peg always makes a decision on which properties it adopts. If it chooses the one set, there may be other pegs from which it is indiscernable that take the other choice, but it can also very well be that no such pegs exist.

Thirdly, it can certainly be the case that our Babylonians have no opinion about the identity or non-identity of Hesperus and Phosphorus, their information state may be one on the basis of which both **may**$(h = p)$ and **may**$(h \neq p)$ is true. They may be forced to add the empirical evidence provided by new astronomical discoveries to their information state, and in that way they will adopt an information state on the basis of which $(h = p)$ is true, there is no reason why they would not be astonished by this evidence. The identity of Hesperus and Phosphorus may be an ontological necessity, but that is an external perspective on their identity conditions. Within the Babylonian perspective the identity of Hesperus and Phosphorus is contingent. The situation is different for an identity statement like $(h = h)$. No information state is able to distinguish the peg Hesperus from itself. $(h = h)$ is certainly true on the basis of the Babylonian information state. Note further that, though identity conditions depend on information states, we can still say that the pegs Hesperus and Phosphorus, *as they are used* by the Babylonians, even if their information state is one on the basis of which Hesperus and Phosphorus are not identical, still refer to one and the same object in the world, the planet Venus. *We* are fully justified to say that the planet Venus was already discovered in Babylonian times, we don't have to say that the Babylonians only discovered two things they pointed at. Even if they believed that the things they were pointing at were the chariots of the evening god and the morning god, respectively, it still is the case that *those things* in fact are the planet Venus. They themselves would accept that, if only they were presented with the right evidence.

Other famous deixis cases, like John Perry, looking in the mirror, saying to himself: "where have I seen that face before?" and David Kaplan,saying while pointing over his shoulder at a picture of Spiro Agnew, which occupies the place of his beloved portrait of Rudolf Carnap: "that is the most important twentieth century philosopher." are analyzed in the same vein.

Let us turn to Kripke's puzzle (Kripke [1979]).

Kripke tells the story of Pierre, a French boy, who during his youth in France hears those fascinating stories about this beautiful town London—in French, of course, so he knows it by the name of 'Londres', but that is no reason for us to deny that on the basis of what he has heard he thinks that London must be beautiful. Then, on a bad day, he is brought to an ugly and dirty suburb of London, to start a new life. He never leaves this suburb and has to learn English the hard way, because none of his neighbours knows even a single word of French. He finds out, of course, that the town he is living in is called 'London', and in days of depression he curses his miserable life, screaming out his sincere belief that London is the ugliest town he knows. Often he sighs in his native tongue: "Que je habitais à jolie Londres!", "If only I lived in beautiful London!"

Kripke's puzzle is interesting, because, unlike most mistaken identity puzzles, the contradiction in his belief does not depend on a clear mistake that he made, on the contrary, he has quite good reasons to believe that the town he was told about is beautiful, and equally good reasons to believe that the town he is living in is ugly (of course, once he recognizes the identity, something has to go). Moreover—as Kripke says—the problem does not make use of dubious principles of substitution, but only of plausible principles of translation (for *us* in describing Pierre's belief).

We can assume that the logical form that Pierre assigns to the sentence-pair **London is ugly. Londres est Jolie.** is the same as we do: $< l_1, U(l_1), l_2, J(l_2) >$. It is very obvious that Pierre's information state is different from ours. We can put the whole problem there, and then there is no difference at all between this case and the Hesperus-Phosphorus case: our information state provides evidence that $(h_1 = h_2)$, Pierre's doesn't.

If we do that, we miss one aspect of the puzzle, namely, that *we* are justified to apply to plausible principles of translation *because* for us h_1 and h_2 are markers for one and the same peg: we can translate because we assume that there is only one object we talk about, whether we talk about it in French, or in English. Obviously, here Pierre's *linguistic* assumptions differ from ours: Pierre clearly talks as if there are two objects. We have to assume that our mapping from discourse syntax to discourse semantics differs from Pierre's mapping. When talking about London, we share the convention that our discourse markers are mapped on the same peg. (This convention can be broken, when two people, after their holydays have

so different experiences of the city of London, that one says: "It seems as if we've been to a completely different town. I don't know about *your* London, but *my* London was beautiful.") But Pierre does not share this convention, his discourse markers are mapped onto distinct pegs.

But we should ask: how can this be? Clearly, when in France he was told about Londres, he and his educators talked about the same peg. Now in England, he and his neighbours also talk about the same peg. But would not his neighbours and his educators agree that there is only one peg *they* talk about (assuming that his educators do know that Londres id London)? And if that is so, how can his French peg be different from his English peg? The explanation is that this takes a more rigid view on pegs than is justified.

Perhaps his neighbours and his educators *would* agree *if they were engaged in a conversation*, but they were not. When we start a conversation, we make some assumptions about the things we talk about, these assumptions belong to the initial conditions which are reset every time we start a new conversation. We agree to talk about London. We know that lots of people have talked about London before, and we know they talked about the same thing. That means that if we talk about their conversations about London, we will use the peg we have agreed upon for that. It does not mean that it is essential for the success of our conversation that we *literally* make sure that we use no other peg than they did, that we use exactly the same thing. We just take the model too serious if we do.

We can say that they talked about the same thing as we do, and they could have said that we would talk about the same thing as they did, and we can say that the peg London plays the same role in our context as the peg London did in their context, but to check whether those objects are literally the same, through contexts which are not part of one conversation, is to take a classical step out of the context. When someone starts a conversation with: "As I said yesterday about London,. . ." the step to conversational success is not to retrieve from your memory the information about which was the exact peg that was used yesterday, but to agree upon a peg now, and ascribe to it the properties that were mentioned yesterday. If we draw the track of London through our increasing information over the years, we may use a different coloured pen from time to time, as long as we make sure that as far as we are concerned it is the same track. When we start a conversation, we agree upon the pegs, like we synchronize our watches. Within the conversation we even assume that they are the same pegs as used in earlier conversation, but we're not going to be bothered about that: if they're not, we identify them. If these French educators of Pierre have set their assumptions on a literally different peg than his English neighbours, then noboby bothers, as long as they are willing to reset these assumptions in a joint conversation.

But that is of course Pierre's problem. He lacks precisely the vital information to reset his French peg in talks to his English neighbours. In that way, he *can* share the same assumptions with his educators *and* with his neighbours. If he would join a French-English dialogue then his problem would soon be discovered, because he will find out that the other participants have reset their pegs, while he hasn't. His discovery then that London is Londres is partly a *linguistic* discovery. It might take place by showing him a French and English map of England.

But apart from that, though he thinks that they don't, in a conversation, his two pegs for London and Londres refer to the same real object, the city of London, and the other participants can show him that they do. In this respect the puzzle is similar to the Hesperus-Phosphorus case.

A third puzzle goes back to Eubulides, Aristotle's famous critic, to whom also the Liar Paradox and the Sorites Paradox are credited. It is called the Paradox of the Hooded Man: You say you know who your brother is. But that man who came in just now with his head covered is your brother. And you said you did not know who he was.

The paradox is discussed as a problem for possible world semantics in Groenendijk and Stokhof [1982]. If my information is represented as the set of epistemic alternatives that I cannot exclude, then I can be said to know who the winner of the 1984-*Tour de France-Feminin* is, if in every alternative that is admitted for me, the one who won the Tour *there* is the one who *in fact* (that is, in the real world) won the Tour. So, to know who is the winner is to have the *specific* information that one and the same individual has won the Tour in all my epistemic alternatives and in fact.

Now consider the following example. We are at a party, and I say to you, pointing at a woman on the other side of the room: "Do you know who that is?" And you answer me: "No, I haven't the faintest

idea." The problem now is that, given the justified trust you have in your perceptual capacities, in all your epistemic alternatives that woman is standing on the other side of the room, and so by definition you know who she is. But yet, you don't. Now I point at a man, next to her, and you say: "I know him alright, he's my brother." For the possible world analysis, the situation is the same, but this time I have no problem in accepting that you indeed know who that is. To pinpoint with somewhat more precision what the problem is, let us make the following additional assumption: at the party there is also a man whom we both cannot tell apart from his twinbrother, although we know them both. Again I ask you who that man is, and you say: "I don't know, I can't tell them apart." Again, I accept what you say as true, without problem: I, as well, cannot tell them apart.

Why do you know who your brother is, and not who the others are? Not because your brother is the same object in all your epistemic alternatives, since the others are as well. You know who your brother is, because you are *acquainted* to him in a way you are not acquainted with the others. In what sense are you acquainted with your brother that you are not with the twins? Both twins are acquaintences of yours. What is crucial for acquaintance as meant here, is that those objects you are acquainted with are the objects that you can tell apart on the basis of your information. But can you not distinguish between all persons at the party, and hence can you not tell them all apart? No, because, though the domain of discussion may be limited to persons present at the party, *we always assume that there are others besides that, who we do not care to distinguish from each other.* The difference is that you are acquainted with your brother in this sense *that you can distinguish him from these others as well.* And that is precisely what you cannot do for the others: you always take into account some, or lots of people outside this room from which you could not tell them apart. For one thing, you may know who has won the 1984-*Tour de France-Feminin*, without being able to distinguish her from the woman at the other side of the room, while you certainly know that your brother didn't.

Being acquainted with your brother in this way, does not imply that you are completely informed about him, acquainted objects need not be total objects: to know who he is you need not know all his properties, it only means that *on the basis of your present information, and in the present context* you can distinguish him from other persons.

We see that if we assume all characters introduced here to be pegs, we can use the definitions of indiscernability and identity on the basis of the information easily to define the notion of acquaintance. But there is a problem here. I said above, that you can be regarded to be acquainted with your brother, even though you don't know all his properties. But then there can always be those nasty pegs that only have properties of which you don't know whether your brother has them or not, and you cannot tell your brother apart from those pegs, so you never know who your brother is. Take the hooded man. Your information tells you only a few things about him: he is a man, he just came in, and he has his head covered. All of them are properties that your information either ascribes to your brother, or about which it is silent. So for all you know this man *could be* your brother. But then your information does not distinguish between the two, and your brother is as much a stranger to you as is the hooded man. And still, when asked, after your brother's identity has been revealed, whether you did know who you brother was and who the hooded man was, you are tempted to say that the hooded man was a stranger for you, but of course in your acquaintance with your brother nothing really has changed. In fact, on thinking about it, this is typically the kind of joke he would enjoy.

The point about acquaintance is not that you can distinguish your brother from any peg whatsoever, because some pegs are so much strangers to you, have so few properties that they can be anyone. Precisely because *they* are complete strangers, you don't take their indiscernability with your acquaintances as an argument that you don't have acquaintances. You don't have to distinguish your brother from all those strangers that you assume to exist, but that you really don't know anything about. *What is crucial about knowing who your brother is, is that you can distinguish him from those pegs about which you know at least as much as you do about your brother.* Your information provides you with a lot of facts about your brother. If all these facts could, as far as your information is concerned, fit different pegs, one of which your brother, then indeed you don't know who your brother is. The basis of acquaintance is not solely the weak notion of indiscernability, which says that two pegs are indiscernable on the basis of the information if they don't have incompatible properties, but it involves the notion of *strong indiscernability* as well: two pegs are

strongly indiscernable on the basis of the information if they have the same properties in it and if the same properties are incompatible with them.

This is how we determine whether we know who your brother is: we compare your brother with all the pegs that are weakly indiscernable from him on the basis of the information, the hooded man being one of them. And then we ask: suppose we add to such a peg all the properties that we know your brother has (we talk about basic properties, of course, and not about logically complex properties like being identical to your brother, in other words, we talk about pegs as if they are pegs and not as if they are classical objects), and similarly, we add to your brother all the properties that this peg has. In other words, suppose we make your brother and this peg strongly indiscernable. Then you are acquainted with your brother if doing that brings you in an information state on the basis of which they are *identical*. You are not acquainted with a peg, if there is a peg indiscernable with it, such that if you make the two strongly indiscernable it is *then* still possible to find new evidence that tells them apart after all.

Take two pegs, of one you only know that he is a young man, of the other you only know that he is a fat man; you can extend your information to a state where you know that both are young and fat men, where they are strongly indiscernable. They are not acquaintances to you, because it is very well possible for your information to grow into a state that tells them apart after all. The same for the hooded man: there can be lots of hooded men, which on some extension of your information are strongly indiscernable from this one, but that in general won't be enough information to regard them as identical. Your brother is a different case: you are so well informed about him that you know that if you gradually find out of a stranger that he has a remarkable resemblance to your brother, there is a stage where this resemblance can no longer be a coincidence, where you find out that he cannot be anyone else but your brother. And if it's not your brother, you will have found that out long before you could have reached the stage where they would have precisely the same properties.

This leads to the folowing definition:

d_1 is **acquainted on the basis of s** iff for all d_2 such that d_1 and d_2 are weakly indiscernable on the basis of **s** it holds that for all $\mathbf{s}' \supseteq \mathbf{s}$: if d_1 and d_2 are strongly indiscernable on the basis of $\mathbf{s}'$ then d_1 and d_2 are identical on the basis of $\mathbf{s}'$

On this definition it is possible to say that you don't know who the hooded man is, but that you do know who your brother is, although you don't have enough information to strictly tell them apart. The same holds for the twin brothers at the party. That case involves three pegs: the twin brothers Bruce and Brian, and the peg you are pointing at. The question asked is: is this Bruce or Brian? Both Bruce and Brian can be acquainted objects for you, there is also no reason why they themselves should be indiscernable: you can very well know that their characters are completely different. But you know far less about the peg standing at the party, he is indiscernable with both of them, and there is no reason to be found in your information to suppose that another peg with the same properties has to be the same object.

Of course, acquaintedness comes in degrees: this peg is far less a stranger to you than is the hooded man. If you know that a peg is strongly indiscernable from the hooded man, then still you have no idea at all who he might be, he may be identified with any man whatsoever. If a peg has the same properties as the peg at the party then the admissible extensions of your information state are narrowed down to four kinds: they can both be identified with Brian, both with Bruce, and one with Brian and the other with Bruce. *Given the present definition of acquaintedness* we can also state the difference between the hooded man and the man at the party without talking about other strongly indiscernable pegs (which sounds more intuitive): the hooded man is highly unacquainted, because there can be lots of acquainted objects it can be identified with. The man at the party is only slightly unacquainted, because we do know that in whatever way our information may grow, it will always be identified with one of the acquainted objects Bruce or Brian. Acquainted objects themselves won't be identified with other acquainted objects than themselves.

Of course, there is a large influence of context (that is, background context) here. You are acquainted with someone only relative to a certain standard of precision. That holds for your brother, it even holds for yourself. You may be acquainted with your brother in the present context, but he may become more of a stranger to you if the context is changed and a different standard of precision is imposed. For instance,

in the present context you can disregard the possibility that, unknown to you, your brother is one of a twin, which in the past changed places every week. If someone, or some evidence, forces you to regard that as a possibility, or even as the fact of the matter, you are forced into a context in which different standards hold: facts which before you did not even have to consider to reject them, suddenly destroy your certainty and you do have to admit to yourself: if I accept those facts, and I reconsider things in this light, then I don't know who my brother is, after all. Dr. Jekyll may change into Mr. Hyde each night without knowing it, he may even hear so much about Mr. Hyde that he can regard himself as acquainted with him. In that context, the question whether he is Mr. Hyde has the same answer for him, as it has for us: of course he is not. Only when faced with the dreadful truth, he has to ask himself seriously who he is.

One might think that the case of the hooded man is a highly constructed one, which we don't often find in real life. Nothing could be more of an illusion than that. In the paper you see a picture of a huge demo which you attended, and you know you must be somewhere in that crowd, but of course you couldn't tell which of those persons is you, you can see the faces of those in the front of the picture, and you know you're not among them, but the rest of them swim in a sea of indiscernability. But it is not *that* that makes you wonder who you really are.

5: Kind hearts and coronets

Kind Hearts and Coronets is a filmcomedy about a man who marries a noblewoman that is eighth in line for an important heritage. Not having patience enough to wait till all these seven relatives have died, he decides to murder them all, and so he does. Towards the end he is arrested for the only murder in the whole movie that he did not commit. A most remarkable aspect of the movie is that all seven noble victims are played by Alec Guinness.

There is something fascinating about that. Suppose you're watching that movie and you see Alec Guinness come in, what do you know about that situation? Well, since Alec Guinness plays all (noble) victims you know that *a victim comes in*. You certainly don't know that all victims come in, even though he plays them all, you even needn't know which one of them comes in. Now suppose that you know that this is the kind of movie where, if Alec Guinness comes in, he wears a hat. What is it that you know about this movie then? Since Alec Guinness plays all victims, you know that *if a victim comes in, he wears a hat*, and that means that *every victim who comes in, wears a hat.*

We see that Alec Guinness is a remarkable actor indeed. If we know that he comes in, we have existential information about the movie. If we know that if he comes in he wears a hat, we suddenly have universal information about the movie.

The sentences

A victim comes in
If a victim comes in he wears a hat

form, of course, part of the Case of the Beaten Donkey(s). I find the analogue persuasive: if we can invite Alec Guinness to live in our model, it seems that we have the right person to be the interpretation of the existential term **a victim**.

The analysis of donkey sentences is the most impressive aspect of discourse representation theory. The elegance of the analysis is most clear in Heim's perspicuous presentation. It comes down to the following ideas:

1. Existential terms are *variables* and not quantifiers.
2. Quantifiers are *nonselective*, they bind every variable in their scope. This amounts to the slight revision of the classical notion of variable binding, by replacing the notion: 'assignment g is identical to assignment f with the possible exception of the value of the variable x' by the notion: 'assignment g agrees with assignment f at least on the values for the variables in set X'.
3. *Existential closure* instead of univeral closure: a formula $\varphi(x_n)$ with free variable x_n is regarded as equivalent to $\exists x_n \varphi(x_n)$ instead of (as usual) to $\forall x_n \varphi(x_n)$: a discourse is true if there is *some* faithful embedding of it in the world (instead of: every embedding in the world is faithful).

4. Conditionals introduce a two-place *universal discourse quantifier*, relating the variables in an antecedent discourse to those in a consequent discourse. This quantifier has the following semantics: for every variable assignment that makes the antecedent discourse true, there is a variable assignment that makes the consequent discourse true and that does not assign different values to the variables of the antecedent discourse.

Given that existential terms are variables, the peculiar behaviour of existential terms in donkey sentences follows immediately: **a victim comes in** has logical form x_1 *is a victim and* x_1 *comes in*, which is interpreted existentially by existential closure; **if a victim comes is he wears a hat** has logical form *(*x_1 *is a victim and* x_1 *comes in)* $\rightarrow$ x_1 *wears a hat)*, where $\rightarrow$ is a universal discourse quantifier that binds x_1.

Pronouns are chameleontic. Sometimes they are very much like proper names, constants, sometimes they are more like variables. A theory that does not want to treat them as ambiguous has to make a choice. Either it reduces constants to variables, while allowing for some rather constant variables, or it reduces variables to constants, claiming that variables are names as well, but—of course—names for rather variable objects. (For an extensive discussion of these alternatives in relation to the semantics of pronouns, see Partee [1984].)

Discourse representation theory is a prime example of the first approach, an example of the second approach is Kit Fine's theory of arbitrary objects (Fine [1984]). I think that the question whether pronouns are variables or constants is not a trivial matter to be settled by taste, but is one of the main intriguing questions in semantics, namely the question how to draw the borderline between syntax and semantics. In fact, it is no other question than the one that was posed before: what is the meaning of a variable?

In what follows, I will choose the side (though not the theory) of Fine in this matter: the meanings of variables are alecs.

6. Alecs

What are alecs?
The answer suggested by the last section is: alecs are actors who play several characters in a play or movie. Of course, real actors are also human beings, like the characters they play: Alecs, similarly, are pegs. In the domain of pegs, we will find alecs and the characters they play.

The hard question then is: what makes a peg into an alec? What properties should pegs have to become alecs? Since pegs don't have properties, this has to mean: what should an information state be like for a peg to be an alec?

We are right in the middle of the movie. We see Alec Guinness come in. He plays the role of a victim. He is not identical to that victim: he plays him, he mimicks his behaviour. At the end of the movie, Alec will have played all victims. One of the victims is already dead, he certainly doesn't come in in this scene, but if Alec has the property of coming in at the basis of our information, Alec will keep that property till the end of the movie, when we have complete information. Then apparently at the end of the movie, we only see Alec in one particular role, apparently, he can't play that role there. Given this perspective, it seems plausible to describe the special nature of Alec as follows:

a is an alec, with respect to being a victim, on the basis of **s** iff for all *d* **s** contains the information that *if d* is a victim, then *d* maybe identical to *a*.

This is nothing but the skolem-function approach to donkey sentences: you associate with *a* for every branch extending **s** the role that *a* plays there. And we know also that this won't work, because it makes alecs existential, also when they are embedded in conditionals.

Another alternative is to impose a partial order of specificity on the domain of pegs, and let the characters be the maximally specific objects, and let Alec be an object that can grow into every victim on every information state; then at the end of the movie, Alec still can grow into every victim, plays all roles.

This is basically Kit Fine's theory of arbitrary objects: Alec is the arbitrary victim. This also won't work, because it makes alecs universal also when they are not embedded: if the arbitrary victim comes in, then every victim comes in.

If we try to identify Alec with one victim in each possible outcome, or with all victims in some, or all

victims in all, then we will unavoidably collaps into one of these alternatives. Also, what do we do with those victims who are already dead? Alec *has* played them, but the present stage of the movie will not be one in which he still can play them, and that is also a factor that will spoil the universal behaviour in conditional contexts, if alecs are existential in non-conditional contexts, and vice versa. The fact that there is a persuasive analogy between existential terms and certain actors does not imply that we should not be very careful in deciding which aspects of the metaphor should be formalized and how: like most analogies, it breaks down at certain points, and it is unadvisable to impose it by brute force.

Let me try to present a slightly different picture.

Let us assume that language users take their present information state $\mathbf{s}$ to be part of a conversational play, which they assume has a beginning point $\mathbf{s}_0$, where $\mathbf{s}_0 \subseteq \mathbf{s}$. $\mathbf{s}_0$ contains the initial conditions of the conversation. The possible courses of the whole play, then are the branches (maximal chains) in the set $\mathcal{B}_{\mathbf{s}}$ of extensions of $\mathbf{s}_0$, the possible end-scenes are the total elements of those branches. If b is a branch, I will use $\mathbf{t}_b$ for its endpoint.

If we want to impose the requirement that Alec plays all roles, then that does not imply that he is identical to any of the roles he plays. I take it to mean that for every character there is a stage in which Alec mimicks the behaviour of that character, there is a stage where Alec's properties and the properties that are added to him there are the same as the properties of that character. We can also say it as follows: there is a stage where Alec and that character are *strongly indiscernable*. However, this cannot be exactly what we want. If we already know of two characters that one wears a red beard and the other a brown beard then Alec can't play both, because each of them will have that beard during all stages, so he can't be strongly indiscernable from both. So Alec does what every actor does if he doesn't have the right outfit: he plays without and lets the audience imagine the beard itself. Alec can't be strongly indiscernable from both, *but he can be an indiscernable approximation of both*: at least the properties that *he* has will be properties that the characters he plays will have as well. Alec has the capacity to *imitate* every character up to a certain stage. Of course, there is always a moment where the imitation becomes obvious, in the end, there is only one real role that Alec can play, namely himself, but as long as we haven't reached that stage, the imitation can be persuasive.

I use this formulation to bring us in the right direction. Taken literally it is far from a solution. For, it has all the problems we could think of for the old alternatives. Take a branch and consider Alec and one of his roles. For some time they act the same. But at the same time, a different role is there that developes in a different direction. How can Alec ever play that role? Should we allow Alec to *change* its role at a branch? But then it has to change its properties at that branch, the properties which according to the information states there it certainly has. If Alec can be non-monotonic in that sense, then he won't be a peg. And in this respect the interpretations of existential terms should not be like actors who change their properties radically in the course of the film, by playing a different role: if we get the information that a victim comes in who wears a hat, we are not going to suppose that we can continue that discourse with: "He doesn't come in", because we mean a different role. That is one point where the metaphor breaks down.

A second point where that is the case is, that if we get the information that a is a victim and a comes in, the requirement that a is an Alec will not mean that he plays the role of all victims (though that can be an initial requirement), but what is relevant for us is, that he plays the role of all victims that come in. In evaluating the existential sentence **a man comes in** on an information state, the property of being an alec will be a three place relation between a peg, the property of being a man that comes in, and an information state. That is, if we say: "A victim comes in", without saying who we mean, we don't have to take into account those victims of which we know that they don't come in.

An alec has to play in some way or other several roles in the play, since at the end of the play he has played them. The fact that this alec has played a certain role, means that between the beginning and the end of the play, he somewhere played that part, somewhere was an indiscernable approximation of that character. Consider a branch b, with beginning point $\mathbf{s}_0$ and endpoint $\mathbf{t}_b$. That branch is the story of the play which ends in $\mathbf{t}_b$, but there are many alternative ways of telling the same story: there are many branches which start in $\mathbf{s}_0$ and end in $\mathbf{t}_b$ (they are not different stories, but different orders in which the same information can be obtained). That an alec plays a particular character in that play means that there

is a particular way of telling that story in which he plays that character, a way of following that alec in a particular role. If we say that in that play he plays different characters, we mean that there are different ways we can follow him through the same play. We can regard those branches as providing the different roles that an alec plays, and if we say that after the play is over he has played several characters, we do not take that to mean that on a particular branch with that endpoint he showed the same behaviour as all those characters, but rather that for each of those characters there is a such a branch where he showed the behaviour of that character.

We will call the branches that have the same beginning and endpoint as a branch b **the alternatives of** b.

This gives us a way of formalizing the notion of an alec as we intend it here:

Alec$(a, \mathcal{P}, \mathbf{s})$, peg a is an **alec** with respect to set of properties $\mathcal{P}$ in $\mathbf{s}$, iff for every branch $b \in \mathcal{B}_{\mathbf{s}}$ and every peg d which at some stage of that branch has the properties in $\mathcal{P}$, there is an alternative b' for b where at some stage a has the properties in $\mathcal{P}$ and a is an indiscernable approximation of d.

The best way of understanding the particularities of this definition, is to see what this notion can do for us in the case of existential terms.

In the course of a conversation someone utters the sentence **a victim comes in**. We change the information state we were in, to a state $\mathbf{s}$ on the basis of which that sentence is true. We do that by choosing a peg a to which $\mathbf{s}$ ascribes the properties of being a victim and coming in, where a is an alec in $\mathbf{s}$ with respect to these properties, i.e.

$alec(a, \{victim, come\ in\}, \mathbf{s})$ and $\mathbf{s} \Vdash victim(a) \wedge come\ in(a)$

So on the basis of $\mathbf{s}$, a is a victim that comes in. The fact that it is an alec tells you that for every branch, extending $\mathbf{s}$ and every object that at some stage of that branch is a victim that comes in, there is an alternative for that branch where a is an indiscernable approximation of it for some time. This gives you existential information: you know that on every way your information can be extended you will find *some* object (namely a) which is a victim that comes in (you do not have to know *which* victim comes in). It does not give you universal information, there is no reason why it would imply that every victim comes in.

Now consider the conditional sentence **If a victim comes in he wears a hat**. To add this to your information you also go to an information state where you have chosen an alec a such that it holds that if a comes in, a wears a hat.

$alec(a, \{victim, come\ in\}, \mathbf{s})$ and $\mathbf{s} \Vdash (victim(a) \wedge come\ in(a)) \rightarrow wear\ a\ hat(a)$

The conditional part of this means, on the definition introduced in the section on data semantics: for every branch extending $\mathbf{s}$ if the first thing that happens with respect to $victim(a) \wedge come\ in(a)$ is that it becomes true there then $wear\ a\ hat(a)$ must be true there. (Note that, in the universal case, a need not, and normally will not, have properties *in* $\mathbf{s}$ itself, even though it is an alec in it. This corresponds to Heim's condition that indefinites introduce new markers. We can say that a is only introduced as an object in $\mathbf{s}$ if it has properties in $\mathbf{s}$.)

Now take any branch where some object d is a victim and comes in somewhere. Then you know that there is an alternative for that branch where a is a victim and comes in and is an indiscernable approximation of d at some stage. But, given the conditional, that means that a must wear a hat there, and since it is an indiscernable approximation of d there, d must wear a hat there as well. This means that every branch where some object d is a victim that comes in, has an alternative where d wears a hat as well. Since the alternatives have the same endpoint as that branch, the fact that d wears a hat will reach that endpoint (which means that in one stage or other it also has to reach that point on the other branch), and we see that on every way our information can be extended with an object that is a victim and comes in, we will find out that it wears a hat as well: in the context of the antecedent of a conditional, the existential term is interpreted universally (see also figure 1).

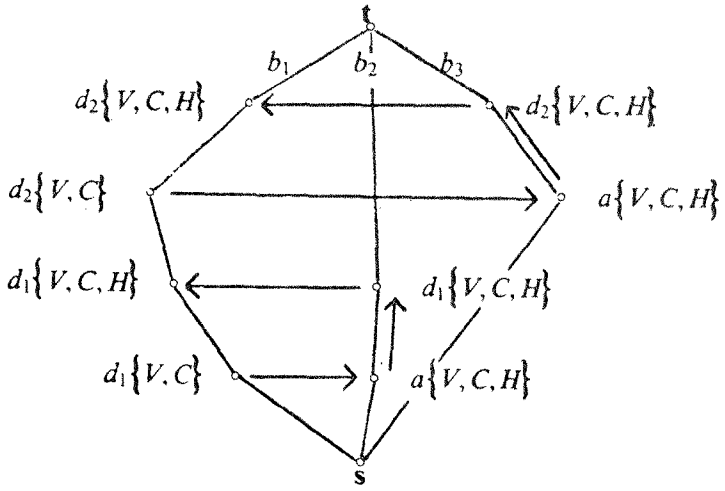

figure 1

These examples use only one existential term. An extension of the notion of an alec with respect to a set of properties, to an alec-pair (n-tuple) with respect to a set of properties and relations is straightforward.

Alecs are pegs, partial objects, which you can follow through your information, about which you can be deadly wrong, which you can mistake for others, and which can turn out to be identical to others. With alecs, it is no mistery any longer how in uttering a hob-nob sentence:

Hob believes that a witch poisoned his cow and Nob believes that she killed his sow

we talk about one object that doesn't exist, that Both Hob and Nob believe to be the same, although none of them believes a particular witch to have done those deeds: because she is an alec, and as an alec she is a peg.

Discourse semantics is the level of objects we talk about, partial objects, that we can learn to distinguish from each other, but objects, that we can also learn to identify with each other. As far as this epistemic contingency is concerned the referents of existential terms are not different from those of other terms: we can discover that they are the same or distinct. This is hard to understand, if we are forced to assume that they are variables, symbols. Fortunately, we don't have to. We can use the partiality of pegs to introduce alecs, objects that express their variability in their behaviour in information growth.

Epilogue

What is the moral of this paper?

It has never failed to be perplexing that people are able to make sense out of the words of others in this jungle of uncertainty, indeterminacy, unspecificity, vagueness, indefiniteness, and indiscernability. Some believe that we can rely on the context to solve all these things for us. Indeed, I would say, the context of shared assumptions of the language users is an important source of efficiency in our information exchange, but not in the sense that it resolves indiscernability. At the contrary. We want to talk about some issues, be it pizza toppings or the meaning of life, but if we had to wait till the context had resolved the vagueness in that, we'd better do something else. But we are impatient. Instead, we let the context set a few pegs for these things, *accept* that we really don't know enough about them, that there can be lots of others, which are more or less indiscernable from them, and start talking about them. The real trick is that we

don't care about our lack of knowledge, that we even exploit it: pegs are the trick that partial information has found to express itself, and indiscernability is the basis of their success, their efficiency. Alecs are so useful because you can let them play the role of practically anything, without having to be over-precise. They would be useless, if our information were complete, maximally precise, etc.: the real world makes too many distinctions to be managable.

Alecs are indefinite objects. But they are *potentially* definite. If you want to talk about something, and you introduce an alec for it, then *by the very nature of an alec* you deliver yourself to the assumption that, how indefinite it may be, how indiscernable on the basis of your present information it may be, you know what it would take in principle to make it definite, *and* you know that in principle it *can* be made definite. You know that *because* you know what it would take for your information to grow and you know that your information can grow. Since pegs and alecs are objects whose indiscernability is determined by your information, the assumption that your information can be made total, and the assumption that you know what it would be if your information were total, has the consequence that they are potentially discernable, potentially specific, potentially definite. Of course, it can be in practice hard to figure out what would make them definite, and that will certainly influence judgements about the acceptability of anaphora in particular cases. But the assumption that we talk about pegs *is* the assumption that exchange of information is about things that, if only we were better informed, we could tell apart. We cannot talk about that of which we have to assume that it is indiscernable in principle. Even if the world were so, that there *are* things that are principally indiscernable, when we talk about them, we introduce pegs for them and therewith pretend that they are not. Where we give up the assumption that we can get better informed, our words lose their intersubjective meaning and are reduced to their form and subjective appearance.

Note

This paper is an abridged version of the paper 'Pegs and alecs' in Landman [1986]. It appears here with permission from Foris Publications, Dordrecht-Cinnaminsion, which is gratefully acknowledged. The research for this paper was financially supported by the Netherlands Organization for the Advancement of Pure Research (Z.W.O.) and by Barbara Partee's grant from the System Development Foundation.

This paper would have been impossible for me to write without the help and stimulation of Emmon Bach, Renate Bartsch, Jeroen Groenendijk, Nirit Kadmon, Barbara Partee, Craige Roberts, Martin Stokhof, Ray Turner, and Frank Veltman.

References

Fine, K.,1984, 'A defense of arbitrary objects', in: Landman and Veltman (eds.), *Varieties of Formal Semantics,* GRASS 3, Foris, Dordrecht

Groenendijk, J. and M. Stokhof,1982, 'Formal and acquainted objects in the theory of information', paper presented at the *Groningen Workshop on Interrogative Quantification.*

Heim, I.,1982, *The Semantics of Definite and Indefinite Noun Phrases,* Diss., UMass., Amherst

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

Kripke, S.,1979, 'A puzzle about belief', in: Margalit (ed.), *Meaning and Use,* Reidel, Dordrecht

Landman, F.,1986, *Towards a Theory of Information. The Status of Partial Objects in Semantics,* GRASS 6, Foris, Dordrecht

Parsons, T.,1981, *Non-existent Objects,* Yale UP, New Yersey

Partee, B.,1984, 'Compositionality', in: Landman and Veltman (eds.), *Varieties of Formal Semantics,* GRASS 3, Foris, Dordrecht

Scott, D.,1980, *Lectures on a Mathematical Theory of Computation,* Oxford University Computer laboratory. Technical Monograph PRG-19