FROM SEMANTICS TO SYNTAX

Martin Kay

December 1967

P-3746

30

FROM SEMANTICS TO SYNTAX

Martin Kay

December 1967

This paper was prepared for inclusion in Recent Developments in Linguistics, edited by Manfred Bierwisch and Karl Heidolf, Mouton, The Hague (forthcoming). An earlier, shorter version was presented at the 10th International Congress of Linguists, Bucharest, 1967.

BLANK PAGE

FROM SEMANTICS TO SYNTAX

Martin Kay*

The RAND Corporation, Santa Monica, California

Language is an instrument of communication. Speakers of a common language exchange tokens and thereby influence what goes on in one another's minds in more or less predictable ways. Two people speak the same language if they not only agree on what counts as an acceptable token, but if they are also able to predict fairly accurately what effect each token will have on the other. It seems, therefore, that a linguistic theory -- a model of human language -- must provide for two kinds of objects corresponding to the tokens and their predictable effects respectively. These can conveniently be labeled sentences and meanings. It must also provide a characterization of the way sentences are paired with meanings.

As Lyons (1963) points out, this is not the only possible view of language. It is true that a token must be meaningful to be linguistic, but this does not entail that there is necessarily something which is its meaning. The so-called structural semanticists view meanings not as entities apart from sentences but as relations among sentences. To say that one sentence is true if and only if

*I am indebted to David G. Hays of The RAND Corporation and Meyer Wolf of Michigan State University for many fruitful discussions of these ideas, but the foolishnesses are mine.

another is true, or that one entails another, or that one is incompatible with another, is to state, at least in part, the meanings of these sentences; it is not to state facts which are true by virtue of relations among some distinct set of entities, the true meanings.

This view is appealing because it avoids one of the traditional philosophical difficulties in discussions of meaning. The meaning of a sentence cannot be separated from its linguistic embodiment; it can be stated only in the terms of some linguistic system, which is to say, in effect, as a sentence. No principle is readily available for chosing one of a set of paraphrases and claiming that it _is_ the meaning of the others. More reasonable seems to be to say that their meaning _is_ the equivalence relation that unites them, and the relations of entailment that sets them off from other such sets.

But, while this view may have some short-term philosophical appeal, it is almost totally barren from the standpoint of scientific linguistics. It leaves unexplained the fact that meanings, as well as relating sentences, do affect people. More important, it fails to explain how people can correctly tell the meanings of new sentences. The number of sentences in a language is usually supposed to be infinite and the number of expressable meanings is presumably not smaller. The fact that all the members of a speech community identify essentially the same set of pairs therefore demands explanation.

It is true that meanings can never be exhibited in a pure form, but only as represented by sentences, but this is not to say that every sentence as semantically as perspicuous as every other. The sentence "The person who lived in the house at the side of theirs bought the car" means substantially the same as "The automobile was sold to their next-door neighbor", but the two sentences have almost no words in common and their grammatical structures are clearly distinct. A semantically perspicuous notation would express similar meanings in similar ways. It would be such that if one sentence entailed another, then the second could be derived from the first by applying certain very general rules.[1] There are no general rules that will derive "One of my parents had an accident today" directly from "My father fell downstairs and hurt himself badly this morning". For a perspicuous notation, the rules could be written.

As I said at the outset, it is part of the work of linguistics to characterize the set of sentence-meaning pairs that make up a language. Since the number of such pairs is infinite, linguistic theory must enumerate them recursively; it will therefore be what is called a generative theory.

[1]This should not be taken as meaning that the particular set of rules required to derive one from the other should itself be discoverable algorithmically. We must expect the semantic formalism to be such that equivalence and entailment relations will be formally undecidable.

Many generative theories are, in principle, possible. Chomsky's theory of transformational grammar has a base component which characterizes a set of base P-markers each of which is translated into a semantic interpretation by a semantic component and into a sentence by transformational and phonological components. A base P-marker therefore corresponds to zero or more sentence-meaning pairs, and the total set generable in this way constitutes the language characterized.

An alternative plan is to make the base, or productive, component, semantic. The base P-markers would then be semantic interpretations, and the transformational and phonological components would translate them into sentences.

There are several objections that can be made to both of these formulations of generative grammar.[2] Most important is the objections that they provide the worst possible basis for an attack on the problems of what Chomsky calls performance. A language is a set of sentence-meaning pairs, but it is also, an instrument of communication. It is therefore of the essence that a speaker of a language should not only know the proper set of sentence-meaning pairs, but that they should be able effectively to translate sentences into meanings and meanings into sentences. He must be able to use his linguistic competence in performance and the linguist must, sooner or later, explain how he does this.

[2] See Hays and Kay (forthcoming).

A linguistic theory which provides a realistic basis for the study of performance must characterize a transducer capable of translating meanings into sentences and _vice versa_. Chomsky's theory certainly does not do this and there is every reason to suppose that it will not be possible, in general, to construct a transducer of the required kind from the rules of a transformational grammar.

Henceforeward, I shall use the term _transductive_ to refer to a linguistic theory which defines a semantic and a phonetic notation, a _transduction rule_ form, an _upward transducer_ or algorithm which translates phonetic expressions into semantic expressions and a _downward transducer_ which translates semantic into phonetic expressions. A particular language will be characterized by a set of transduction rules which specify the operation of both transducers.

The grammar of a language must distinguish a set of well-formed phonetic expressions, or _sentences_, and a set of well-formed semantic expressions, or _meanings_. The question of whether the set of meanings is a linguistic universal is important, but it would be premature to try to solve it now. Another important question, which must be answered at least provisionally before any further progress can be made, concerns the ranges and domains of the transducers and the way in which sets of well-formed expressions are to be defined. The simplest solution is to require that the transducers produce at least one well-formed expression as output in response to every well-formed input and that they produce no output in response to an ill-formed input.

A second possibility is to require that the transducers produce well-formed output in response to every input. In such a system, the well-formedness of an expression could be decided by translating it into one or more expressions of the other type with one transducer and translating each of these into one or more expressions of the original type with the other transducer. The expression is well formed if, and only if, it is found among the expressions resulting from the second transduction.

This plan is attractive because it offers the possibility of explaining how ungrammatical sentences are understood and how grammatical sentences can be only partially coherent, or cognitively anomalous.

A host of other possibilities suggest themselves. We can imagine a system which mapped a sentence onto meanings and non-meanings according to the various natural and anomalous interpretations it had and which mapped non-meanings sometimes onto sentences and sometimes onto non-sentences. In this case, however, the transducers themselves make no distinction between well- and ill-formed expressions and the grammar would therefore have to contain at least one extra component with a function similar to that of Chomsky's base component. But if this were found necessary in a satisfactory theory, it would weaken the claims that could be made for the transductive theory because these rest heavily on the proposition that a base component cannot be satisfactorily motivated. At this stage, it is altogether unclear which approach will be most fruitful.

The theory to be sketched shortly takes the simplest approach. Well-formed expressions are carried onto well-formed expressions and ill-formed expressions give no output.

In what follows, I shall describe a form of rules which translate semantic expressions in a notation akin to that of the predicate calculus into P-markers. The same rules, interpreted by a slightly different transducer will translate P-markers into semantic expressions. I have chosen to concentrate on the downward transduction process partly because the operation of the upward transducer is simpler and partly to stress the difference between the view of language advocated here and the more fashionable one in which semantics is regarded as purely interpretive.

To say that we propose to write semantic expressions in a notation akin to that of the predicate calculus is, on the face of it, to say very little because the notation is rich enough to talk about objects of any kind and to describe any kind of relational structure over them. There is no limit to the number of ways in which even so simple a sentence as "The man ran" could be translated. The following are some possible candidates:

```
ran(man)                                              (1)
run(man) past(run)                                    (2)
run(human) male(human) past(run)                      (3)
def(human) male(human) run(human) past(run)           (4)
```

In a sentence consisting of a subject and an intransitive verb, we are used to thinking of the very as specifying properties predicated of the subject. Example (1) therefore seems to be a natural translation. We know however that, for many purposes, the feature of tense needs to be separated out, and this is done in (2). Here "run" must be thought of as implying no particular temporal reference. The past tense is represented by conjoining a second elementary proposition in which "past" is predicated of "run". There is reason to believe that it will be profitable to separate out the feature of sex as is done in (3). The content of the English word "man" is now represented by a pair of primitives, namely "human" and "male". Example (4) is an attempt to accommodate the definite article.

It is by no means clear at what point the process of factoring notions into simpler notions should stop. Outside logic itself it is difficult to find a notion so elementary that it cannot, with a little thought, be expressed as a combination of others. But this is not to say that the entire enterprise is misbegotten. It is indeed possible to state properties that semantic expressions with any real explanatory value must have and to develop a theory of the lexical structure of a language by considering examples and counter-examples and the semantic expressions that can be associated with them by a set of formal rules.

As I have said, semantic expressions must be perspicuous. If two sentences are judged to be exact paraphrases of one another, then their semantic representations must be equivalent. If one sentence logically entails another, then the semantic representation of the second should be derivable by formal rules from that of the first. This is a reasonable requirement of a semantic theory because a notion of entailment is, in fact, often crucial to deciding whether a text is well formed. Consider a fairy story which begins as follows:

> "Once upon a time a lion and a unicorn decided to go hunting together. But one of the animals was afraid of the other because of what he might do when they got into the forest. It was the lion that was afraid of the unicorn."

If the second sentence is left out, the result is no longer acceptable as the beginning of a well-behaved fairy story. The sentence "It was the lion that was afraid of the unicorn" can be introduced only when it has been established that a lion and a unicorn are involved in the story, and that someone is afraid of someone else. Notice however that the following is entirely acceptable:

> "Once upon a time a lion and a unicorn decided to go hunting together. The lion was afraid of the unicorn."

The sentence "The lion was afraid of the unicorn" is true if and only if the sentence "It was the lion that was afraid of the unicorn" is true. But the latter is unacceptable in at least some contexts where the former sounds entirely natural. But for the latter to be acceptable there must be some previous sentence in the text which implies that somebody is afraid of somebody. A semantic theory must therefore provide some formal mechanism for deducing "somebody is afraid of somebody" from "but one of the animals was afraid of the other."

Another story might begin as follows:

> "John Doe was a school master who, at the age of seventy, fell in love with an airline hostess. The girl decided to give up her job and marry him."

If not in the best literary style, this is at least more acceptable than the following:

> "John Doe was a school master who, at the age of seventy, fell in love with an airline hostess. The boy decided to give up her job and marry him."

Second example sounds strange for several reasons, one of which is that no boy has been introduced into the story. There is no parallel difficulty in the first example because airline hostesses are, in general, girls. It is presumably up to the semantic component of the grammar to record this fact and to infer that one of the characters in the story is a girl because one of them is an airline hostess. The sentence "John Doe was a school master who, at the age of seventy, fell in love with an airline hostess" presumably differs in its semantic representation from "John Doe was a school master who, at the age of seventy, fell in love with a girl" in containing a few extra predicates -- those necessary to make the meaning of "girl" more specific so that it becomes the meaning of "airline hostess".

It is not the purpose of this paper to provide even the most tentative set of semantic primitives. I am concerned only with providing the formal system which will relate expressions involving semantic primitives with sentences in a real language. This is, of course, a much simpler task, but it must precede any serious work on the primitives themselves, because, in the absence of rigorous method of associating sentences with semantic representations, there is no sure means of judging the effect of introducing or deleting this or that primitive, and the discussion rapidly loses all contact with reality. However we cannot proceed to set up the formalism without taking notice of some more substantive issues and the most impottant of these concerns the relationship between the expressions we shall write and expressions in a genuine logic such as the first-order predicate calculus.

The humanist is quite right to claim that language is not logical if, by that, he means that the language of semantic representations cannot itself be a logic. Expressions in natural language can be logically equivalent without being linguistically equivalent. The sentence "It was the lion that was afraid of the unicorn" is true if and only if the sentence "The lion was afraid of the unicorn" is true. But, as we have seen, the latter can be used in places where the former is anomalous. It is, of course, possible to claim that pairs of sentences such as these should be regarded as having the same meaning and that the differences between them should be accounted for on some other basis. The grammar might, for example, contain a kind of super-syntactic component describing the way in which sentences are collected to form discourses. Or there might be a stylistic or rhetorical component, concerned with such matters as emphasis, in which these sentences would be distinguished. But, for reasons I shall present shortly, I prefer to regard these sentences as having different meanings even though they are logically equivalent.

The second reason why the language of semantic representations cannot itself be a logic is that it must be capable of expressing certain notions which must be regarded as meaningless in any logic. There must, for example, be no restriction on the ways in which predicates can be applied to other predicates and, in particular, there must be no restriction against a predicate being applied to itself. Allowing this is, of course, exactly what leads to Russel's famous paradox for, if of any one place predicate p the proposition p(p) is meaningful, then we can construct the predicate q such that q(p) is true if and only if p(p) is false. But, substituting q for p, we have the logical fact that q(q) is true if and only if q(q) is false.

This paradox has led logicians to reject any system in which predicates can be applied indiscriminately to other predicates and to deny the meaningfulness of such notions as self-predicability. But self-predicability is not meaningless. The word "short" is itself short and is therefore self-predicable; the word "long" is not long and is therefore not self-predicable.

It is probably true to say that Russell's paradox came to light precisely because natural languages are capable of expressing ideas which logic cannot accommodate. If such terms as "self-predicable" and "the set of all sets" did not sound beguilingly natural, they would never have caused any difficulty. And if they were not meaningful then it is difficult to see how they could have been so closely analyzed as to show their underlying illogicality. The point is this: Logic must be purged of contradictions if it is to be useful for its intended purpose, namely to construct arguments that can be faulted only be faulting the premises on which they are based. Such a system is just as necessary for constructing arguments about language as anywhere else. But our concern here is not with constructing a metalanguage for linguistics, but with constructing a language capable of expressing the things that people express in normal speech. It is true that we have said that the possibility of making inferences in this language by means of formal rules will be one of its crucial properties. But these inferences will have their basis in the culture and the structure of the language, and not in axioms and theorems. The examples (1)-(4) are not therefore to be rejected simply because the same word appears now as a predicate and now as an argument. But they are unsatisfactory from at least one other point of view, namely that they fail to make a necessary distinction between intention and extention. Translating the sentence "The boy saw the man" in a similar way, we might have something like:

see(human, human).male(human).male(human).
young(human).past(see)

assuming, perhaps unrealistically, that the representation of "boy" differs from that of "man" only by the additional predicate "young".

This is clearly unsatisfactory because each instance of the term "human" refers to one of two different people and there is no way of telling them apart. Another possibility is the following:

human(x).human(y).male(x).male(y).young(x).
see(x,y).past(see)

Here the variables, x and y, correspond to referents. Reference is another notion that needs to be interpreted more broadly in a linguistic than a logical or philosophical context. It is clearly not profitable to think of referents as objects in the real world because it is highly unlikely that the sentence "The boy saw the man" in which "the boy" and "the man" may refer to real people is interpreted in a different way from "The unicorn saw the lion" in which the unicorn is presumably fictional. In fact, from a linguistic point of view, there is reason to suppose that the sentence like "The boy saw the man" involves more than two referents because there are also sentences like "The boy saw the man and when he did he ran away". The word "did" in this sentence clearly refers to the same event as the word "saw". In the sentence "The boy saw the man and so did the girl" on the other hand, the word "did" refers to a different event, though one which can have the same name.

The rules for developing P-markers will usually contain two parts, one referring to conditions that must hold in the semantic representation if the rule is applied and the other referring to conditions that must hold in the P-marker, as so far developed, and specifying additions to be made to the P-marker. The second part of a rule will look essentially like a sequence of one or more context-free phrase-structure rules of the familiar type. Let us assume for the moment that the representation of the sentence "The boy saw the man" is as follows:

human(x).human(y).male(x).male(y).young(x).
see(z).past(z).z(x,y)

Some of the rules that might be used in producing the P-marker from this expression are as follows:

(i) a(b,c) | $S \rightarrow NP_b\ VP$, $VP \rightarrow V_a\ NP_c$

(ii) see(a).past(a) | $V_a \rightarrow$ saw

(iii) human(a).male(a).young(a) | $N_a \rightarrow$ boy

The two parts of each rule are separated by an upright bar. Variables (a, b, c...) appearing in the first part of a rule can match any variables (x, y, z...) in the semantic expression. Rule (i) is interpreted as follows: If the semantic expression contains one variable predicated of two others and if the symbol "S" appears at some node in the P-marker which does not, so far, have any descendents, then the symbols "NP" and "VP" can be written against two new nodes dominated by the "S". Two more new nodes are to be introduced beneath "VP" labelled "V" and "NP" respectively.

The subscripts appended to certain symbols in the second half of a rule allow nodes in the P-marker to be associated with variables in the semantic expression. Thus, in rule (i), the "NP" node introduced below "S" is to be associated with the first argument of the predicate translated by this rule. Similarly, the newly introduced node "V" is to be associated with the predicate itself and the second "NP" is to be associated with the second argument of the Predicate. Rule (ii) is to be interpreted as follows: If "see" and "past" are predicated of the same variable in the expression and the P-marker contains a node labelled "V" which has been associated with that variable, and if, furthermore, that node in the P-marker has no descendent, then a node labelled "saw" may be introduced beneath it. Interpretation of rule (iii) is similar. The "boy" can be introduced into the P-marker beneath a node labelled "N" provided that this has been associated by some previous rule with a variable in the semantic expression of which "human", "male" and "young" are all predicated.

A fourth rule might be as follows:

(iv) human(a).male(a) $|$ $N_a \longrightarrow$ man

We are assuming, perhaps unrealistically, that the representation of "boy" differs from that of "man" only by the addition of the predicate "young". It apparently follows that wherever rule (iii) applies rule (iv) will apply also. In principle it would be possible to avoid conflicts of this sort so that whenever the semantic properties specified by one rule included those specified by another, then the more inclusive rule would always be applied first. But this solution will not work. In the first place, since one of the main aims is to construct a model which will account for paraphrase, it must be possible to derive several different P-markers from the same semantic expression. This can only be achieved if rules are in general optional and there is little to be gained by ordering optional rules. In the present case nothing important is lost because if our assumptions about the semantic representations of "boy" and "man" were correct, then it would presumably be at least in part because we were prepared to accept any sentence containing the phrase "young man" as a paraphrase of the same sentence with this phrase replaced by "boy". The following would presumably be included in the set of rules:

$$\text{(v)}\quad \text{young(a)} \;\Big|\; \text{ADJ}_a \longrightarrow \text{young}$$

Using this rule we can produce the sentence "The young man hit the man" from the same original semantic expression.

As the rules are refined to represent more and more semantic facts about the language, the expressions that will have to be written to represent even very simple sentences are likely to become long and involved. Furthermore, since the rules are optional and since some may be applicable at many places the same expression, the process of producing a P-marker may often block, i.e., reach a place where it cannot continue owing to lack of applicable rules, but where some parts of the expression have still not been translated. We can imagine the strategy which continues to explore different sequences of rules until it finds one which yields a translation of the expression or demonstrates that no such sequence can be found. But this is both tedious and unrealistic. If a grammar contains the rule (i) suggested above, then any two-place predicate could be used as the starting point for a P-marker. The same would presumably also have to be said of one place predicates. Hence they will be the normal representatives of intransitive verbs. The production of a sentence could therefore begin almost anywhere.

However, expressions that adequately represent the meanings of sentences will almost certainly have to contain rhetorical predicates, that is, predicates that do not influence the truth value of the sentences in which they occur but which relate the information in a given sentence to a whole discourse--in other words, which present the thread of the argument.

The sentence "It was the lion that was afraid of the unicorn" differs from "It was the unicorn the lion was afraid of" in the information that must be assumed for its understanding and the information it purports to offer as new. In the first sentence, it is taken as given that someone was afraid of the unicorn; the question answered by the sentence is "Who?". In the second sentence it is given that the lion was afraid of someone, but who was he afraid of? The distinction comes out particularly clearly when the sentences are embedded in certain kinds of other sentences. Compare, for example, "I do not believe that it was the lion that was afraid of the unicorn" with "I do not believe it was the unicorn the lion was afraid of". What I disbelieve is very different in the two cases. In the first case, I positively do believe that the unicorn was feared, but not necessarily by the lion; in the second, I believe that the lion was afraid, but not necessarily of the unicorn.

Consider now the sentence "I do not believe that the lion was afraid of the unicorn". Here, what I disbelieve is either that the lion was not afraid or that the unicorn was not feared, or both. In other words, this sentence is ambiguous; it can have the meanings of either or both of the other two.

In the light of this, let us introduce the two-place rhetorical predicate "new" and represent "It was the boy that saw the man" somewhat as follows:

human(x).human(y).male(x).male(y).young(x)

see(z).z(x,y).past(z).new(x,z)

"new(a,b)" means that the predication "b(...,a,...)", which must appear in the same conjunct, gives new information. The first rule used to generate the sentence "The boy saw the man" could now be

$$a(b,c).new(b,a) \left| \begin{array}{l} S \longrightarrow NP_b \; VP \\ VP \longrightarrow V_a \;\; NP_c \end{array} \right.$$

However, since no transformations are to be included in the grammar so that the output of this transducer will have the same status as surface structure in transformational grammar, the VP node created by this rule has no obvious function outside phonology. But for phonology, the rule could be written

$$a(b,c).new(b,a) \mid S \longrightarrow NP_b \; V_a \; NP_c$$

The cleft sentence "It was the boy that saw the man" is also generable from the expression (5). In this case, the initial rule must be something like

$$a(b,c).new(b,a) \quad S \longrightarrow it \; BE_a \; NP_b \; that \; V_a \; NP_c \qquad (5)$$

If the extra nodes are needed for the phonology, the rule would have to be something like

$$a(b,c).new(b,a) \quad \begin{array}{l} S \longrightarrow NP \; VP \\ NP \longrightarrow N \\ N \longrightarrow it \\ VP \longrightarrow V \; NP \\ V \longrightarrow BE_a \\ NP \longrightarrow NP_b S \\ S \longrightarrow NP \; VP \\ NP \longrightarrow N \\ N \longrightarrow that \\ VP \longrightarrow V_a \; NP_c \end{array}$$

A slight variation of the formalism would allow

$$a(b,c).new(b,a) \;\Big|\; S \longrightarrow NP(N(it))\ VP(V(BE_a)\ NP(NP_b\ S(NP(N(that))\ VP(V_a NP_c))))$$

In what follows, I shall ignore phonology and assume that the simpler form of rule can be used. These can always be expanded into the more complex form if necessary.

The sentence "It was the man the boy saw" is derived from a semantic expression including "a(b,c).new (c,a)" and the initial rule is

$$a(b,c).new(c,a) \;\Big|\; S \longrightarrow it\ BE_a\ NP_c\ (that)\ NP_b\ V_a \quad (6)$$

The similarity between (5) and (6) can be captured if they are both replaced by

$$a(b,c).new(d,a) \;\Big|\; S \longrightarrow it\ BE_a\ NP_d\ that\ C_a$$

$$a(b,c).new(b,a) \;\Big|\; C_a \longrightarrow V_a\ NP_c$$

$$a(b,c).new(c,a) \;\Big|\; C_a \longrightarrow NP_b\ V_a$$

The symbol "BE" will be replaced in a subsequent rule by an appropriate part of the verb "to be". The subscript will allow the proper tense to be chosen.

A distinction between given and new, topic and comment, theme and rheme has been discovered and rediscovered from time to time throughout the history of linguistics. In many languages, like Chinese and Japanese, it is reflected in the gross syntactic structure of sentences and it must therefore by explained to beginning students. It has been a major center of attention for the Prague school because, while there is relatively free order among the major constituents of sentences in the slavic languages, the diverse possible orders have different rhetorical effects. It is doubtful whether only one rhetorical predicate, the one I have called "new" underlies all these distinctions. Certainly, there are others and the kind of linguistic theory advocated here would provide a very apt framework in which to study them.

The principal motivation provided so far for the predicate "new" in this discussion, namely of marking possible starting points for the transducer, is, in a sense, a very old idea. Every semantic expression that underlies a declarative sentence must have at least one occurrence of the predicate "new". But this is only to recall one of the most ancient definitions of what a sentence is, namely that it introduces a new thought.

I have already remarked that, in sentences like

"I believe it was Brutus that killed Caesar" (7)

and "I believe it was Caesar that Brutus killed" (8)

the embedded sentences are, by the usual criteria, logically equivalent but that what I am claiming to believe is different. Using the notions of given and new, (7) and (8) can be paraphrased somewhat as follows:

"I believe (given that Caesar was killed, Brutus was responsible) (9) and "I believe (given that Brutus killed someone, Caesar was the victim) (10)

Now, notice that what is given can be removed from the parentheses without change of meaning, whereas what is new cannot. (11) and (12) are not semantically distinguishable from (9) and (10).

"Given that Caesar was killed, I believe (Brutus was responsible) (11)

"Given that Brutus killed someone, I believe (Caesar was the victim) (12)

This suggests liking rhetorical predicates with what Ross and others call <u>performatives</u> and what the philosopher Austin calls <u>illocutionary acts</u>. Ross suggests that declarative sentences come from an underlying structure in which the highest-level sentence is something like "I assert that S". Austin points out that, as well as enshrining facts about the world, sentences can also change those facts. This is particularly true of sentences containing words like "promise", "undertake" and "assert".

If we say that every sentence involves at least one performative in its semantic structure, or that every sentence is an illocutionary act, we can represent (13) and (14) as (15) and (16).

"It was Brutus that killed Caesar" (13)

"It was Caesar that Brutus killed" (14)

Brutus(b).Caesar(c).killed(x,c).b=x (15)

Brutus(b).Caesar(c).killed(b,x).c=x (16)

Clauses like "b=x" now take the place of the "new" predicate. The representation of sentences like (7) and (8) becomes something like (17) and (18).

Brutus(b).Caesar(c).I(i).killed(x,c).believe(i,b=x) (17)

Brutus(b).Caesar(c).I(i).killed(b,x).believe(i,c=x) (18)

Rhetorical predicates, or devices like the "=" in the above examples are crucial to the study of semantics in natural languages and serve to set those languages off against most artificial languages. Just how many of them there are is, of course, not known -- probably very few. The "new" predicate is certainly responsible for the distinction between restrictive and non-restrictive relatives, the latter being precisely those that involve new information, or, more precisely, information that purports to be new. It may also underlie the distinction between "The house is near the store" and "the store is near the house". It almost certainly does not account for the difference between "the glass is half full" and "the glass is half empty", or between "who am I?" and "which is me?".

Informally, the function of a rhetorical predicate is fairly clear. When I talk to you, I normally have a construct in my mind and I want to erect a copy of it in yours. I must build the picture piece by piece, telling you where each new piece is to be placed relative to what has already been built, and checking from time to time that the construction is proceeding as I intend. Each piece comes in a sentence which also contains the necessary information about where it is to be put. It must, therefore, contain old and new information to be effective -- probably more old than new. What is new does not always have to be marked explicitly because it will be clear from what has been done already and the course that the construction is taking. A sentence like "It was Brutus that killed Caesar" marks what is to be taken as new very clearly, whereas the simpler sentence "Brutus killed Caesar" is relatively ambiguous.

But it would be absurd, at this stage, to make any psychological claims for the model of language sketched here. Far too many details have been left out. But, though the specific proposal made here may turn out wrong in almost every detail, and though there is no reason to claim any psychological reality for it, there is nevertheless very great importance in having models in linguistics which actually work -- models of performance or, at least, performance-oriented competence. A model which claims to account for paraphrase must produce paraphrases even if some of the procedures it uses are of no eventual theoretical interest. The trouble with a purely interpretive theory of semantics is that, though one may argue about it endlessly, it is very difficult to see what it would be like to demonstrate that a grammar written according to it was wrong. In the interest of intellectual hygiene alone, it is important to study a subject as difficult as semantics with models that actually work and which therefore lay themselves open to specific attack.

REFERENCES

1. Hays, D. G. and M. Kay, The Failure of Chomskian Theory in Linguistics, The RAND Corporation, (forthcoming).

2. Lyons, J., Structural Semantics: An Analysis of Part of the Vocabulary of Plato, Oxford, Blackwell, 1963.