REVIEW ARTICLE

Stanley Peters & Esa Saarinen (eds.), *Processes, Beliefs and Questions*, Synthese Language Library Vol. 16, D. Reidel Publishing Company, Dordrecht, Boston and London, 1982, XXXI + 231 pp.

Reviewed by JEROEN GROENENDIJK & MARTIN STOKHOF
University of Amsterdam

The last decade has seen the rise of formal, model-theoretic semantics in linguistics. Receiving an impetus both from philosophy and logic, as well as from transformational linguistics, semantics has grown into an interdisciplinary field of study that has brought us descriptive successes and theoretical insights that are quite remarkable in view of its short history.

Yet already some see the dark shadows of a foundational crisis looming. The not always clearly noticed fundamentally anti-psychologistic and anti-mentalistic background of the philosophical tradition from which modern-day semantics springs, seems to become a major obstacle in two areas of research, that of lexical semantics and that of the semantics of propositional attitude verbs. This holds specially for intensional, possible worlds semantics, which traces its origin back to the works of Frege and acquired its present form from the work of Hintikka, Kripke, Montague and others. Within the frame of reference of model-theoretic semantics alternative approaches are being developed which are partly motivated by the problems mentioned above. Others have been attracted by work done in fields which hitherto have not been in close contact with linguistics and philosophical semantics, viz. cognitive psychology and artificial intelligence. Kamp's work on discourse representation, for example, is meant to 'fuse' modeltheoretic semantics and insights from procedural semantics as it is being developed in work on artificial intelligence.

In view of the character of the two main problem areas that semantics faces today, an exchange of ideas and perhaps even a closer cooperation between linguistics, philosophy and psychology in the area of semantics

Studies in Language 9:3 (1985), 409–427. DOI 10.1075/sl.9.3.07gro
ISSN 0378–4177 / E-ISSN 1569–9978

seems a viable option. For the topics indicated have a common characteristic: they all have to do with information. In lexical semantics the everlasting problem is how to distinguish between 'pure' aspects of meaning and knowledge of the world. The work of Putnam and others has revealed a sharp contrast between intensions, the intensional semanticist's reconstruction of the meaning of a word, and what can reasonably be said to belong to the semantic competence of actual language users. Knowledge and prejudice, expertise within a language community and merely prototypical information on the part of the individual seem to play a role in creating the meaning of a word. As for propositional attitude verbs, they are just that: expressions that refer to attitudes of language users. Prominent among these are the epistemic attitudes: belief and knowledge. Also one might point here to other classes of expressions such as modals, constructions such as indicative and counterfactual conditionals, and the like which, some claim, are concerned with information as well. And of course there is that field of meaning which is 'pragmatic' in the sense that it pertains to conditions governing (correct) language use. Gricean theory has revealed a lot of non-truthconditional, but systematic aspects of meaning that need to be covered and that contain essential reference to epistemic states of language users.

So from within, as well as from outside, there is pressure on semantics to question its 'splendid isolation' from all psychologistic and mentalistic talk, and to turn to developments in psychology and artificial intelligence. And, on their part, these disciplines may come to show a broader interest in formal semantics, once it is realized that a theory of whatever goes on in the mind ultimately needs to account for the fact that the human mind, at least for a large part, is orientated outwards, towards the world, and that the human language is shaped by this. Also, specifically for artificial intelligence purposes, the formal apparatus of modern semantics and its associated standards of explicitness and formal rigor may become increasingly important.

Of course a rapprochment of disciplines is no easy matter. Many questions, theoretical and practical, have to be answered and many positions have to be thought over and reconsidered. In view of this the essays collected in this book are very welcome indeed. The papers, by well-known semanticists such as Barbara Partee, Max Cresswell and Nuel Belnap and equally well-reputed psychologists and (psycho)linguists such as P.N. Johnson-Laird, Janet Dean Fodor, Robert Moore, Gary Hendrix, and Willem Levelt, either address the methodological issues explicitly, or, by giving case studies, provide indirect support for one of the positions one may hold in this area. In two

contributions, those of Cresswell and Belnap, a plea may be found for semantics as an autonomous discipline, whereas the papers by Fodor and Moore & Hendrix stress the autonomy of psychology. Partee and Johnson-Laird each in their own way argue for a combining of forces. They deal with the matter of the relationship between formal semantics and psychology explicitly, Johnson-Laird discussing many separate topics, Partee restricting herself to the problem of the semantics of belief-sentences. Moore & Hendrix and Cresswell discuss the latter topic too, each from a different perspective. Belnap addresses the main theme briefly but outspokenly, and devotes the remainder of his paper to a detailed study of yet another topic in semantics that has to do with information, that of questions and answers. Fodor's paper is a case study too, on the representation of quantifiers, and her remarks about other treatments concern mainly linguistic proposals in the Chomskyan tradition. From Levelt's paper no view on the main issue can be inferred. His is a study of a specific psycho-linguistic problem, that of linearization in discourse.

The book contains an excellent introduction by the editors Esa Saarinen and Stanley Peters, who give a lucid sketch of the various issues involved, and excellent summaries of the individual papers. We will now briefly review the individual contributions.

Johnson-Laird's lengthy paper 'Formal Semantics and Psychology' starts out with the observation that logicians have only studied the relationship between language and the world and that psychologists have only related language to the mind. The real task, however, he says, is "to show how language relates to the world through the agency of the mind". His paper is meant to contribute to this by bringing together formal and psychological semantics and investigating what each discipline may learn from the other. Consequently, Johnson-Laird devotes quite a large part of this paper to a characterization of the principles underlying the two enterprises and to an exposition of various positions and results.

The features he ascribes to formal semantics are those which characterize intensional semantics: the notion of an intension as a function from possible worlds and moments of time to extensions; the principle of compositionality of intensions; and the intensions of basic expressions being taken as primitive. Johnson-Laird takes as common to psychological theories of meaning the principle that meanings are represented by expressions of an internal mental language, so-called 'propositional representations'. Propositional representations are assumed to be sufficient to account for various

semantic properties and relations in natural language. Johnson-Laird argues that this is not correct and that psychologic semantic theories are incomplete. It is necessary to take into account the way in which sentences are mentally related to the world. Language users build mental models of (parts of) the world. These mental models correspond to a different, deeper level of understanding than the propositional representations.

Johnson-Laird then considers the question of whether perhaps mental models may serve to link up formal and psychological semantics with a neat division of labour on the side: formal semantics describing what is being computed, psychological semantics how the computation is done. Johnson-Laird notes four major potential problems: (i) lexical semantics: the formal semanticist's intensions cannot be interpreted psychologically; (ii) representation of the infinitude of possible worlds; (iii) completeness: formal semantics uses complete models whereas mental models are characteristically incomplete; (iv) propositional attitudes. Each of these problems is discussed by Johnson-Laird separately.

In formal semantics each basic expression is assigned an intension, a set of necessary and sufficient conditions for something to be in the extension of the term. Such intensions are not plausible candidates for what a language user actually has at his disposal, as Putnam has shown for natural kind terms. For many terms no 'real' intensions are available, and for those for which there are, it is clear that we can never be sure whether we have grasped them. Also, there is no reason to suppose that extensions of terms are always fixed, that fixing the extension is part of learning the meaning of a term. So real intensions need not worry the semanticist. Such arguments, equally fatal for propositional representations as Johnson-Laird notes, do not show, however, that a coherent semantic theory is impossible. From psychological evidence Johnson-Laird concludes that stereotypes can be regarded as effective intensions (in a non-technical sense) which can be given a mental interpretation. He argues that they may take the form of a Minskyan frame, a structured piece of knowledge containing variables which may be assigned default values: values which they are assumed to have, unless there is concrete evidence in a situation to the contrary. Associating such stereotypes with lexical items, clearly makes their meaning knowable, and thus removes one obstacle in the project of linking formal and psychological semantics.

The second problem Johnson-Laird addresses is that of the possible psychological reality of possible worlds. He proposes a constructivist approach to their mental representation. People actively construct mental models which are alternatives to a given situation. But these alternatives are

few in number and depart from the original situation as little as possible. Although such a constructive approach seems to square with certain model-theoretic analyses, Johnson-Laird concludes that in this respect the relation between psychological performance and model-theoretic semantics remains rather remote.

The third problem which arises in an identification of mental models with the models of formal semantics, concerns the completeness of the latter, which contrasts with the incompleteness of the former. Mental models are always radically incomplete since no discourse provides enough information to build a complete one. Hence, in contrast to the formal semanticist's models, mental models will leave the truth value of many statements undecided. One way to represent partial information is in the form of alternatives. Johnson-Laird argues that this is not a viable option: the number of alternatives in most cases would be far too large to handle. He proposes therefore to augment mental models with other representational devices, such as propositional representations. Such features of representation are lacking in formal models. But it seems, Johnson-Laird argues, that they could be introduced, and hence (in)completeness would not have to be a real obstacle.

The last problem, or rather cluster of problems that Johnson-Laird addresses concerns propositional attitudes. He sketches an intensional approach, in the vein of Hintikka, and concludes that it fares rather well, but leaves two problems open: the status of possible individuals, and the problem of substitution of equivalents.

As for the way possible individuals are treated in formal semantics, Johnson-Laird objects to the rigid designator view of proper names. The consequent necessity of identity statements involving proper names, he says, "strikes a psychologist as decidedly odd". A constructive approach need not run into similar troubles. Since the construction of a mental model of an alternative situation proceeds from the actual situation, problems as to which individuals one may or should postulate and how they are to be identified do not arise. (This line of thought, by the way, seems much closer to the way Kripke thinks of possible worlds than Johnson-Laird seems to realize.)

The problem of equivalent expressions is that they are assigned the same intension where obviously some such expressions may differ in meaning. Various solutions which have been proposed are discussed and rejected, on the grounds that they lead to over-refinement and yet still leave some problems unsolved. Rather, as Johnson-Laird puts it, "the idealizations of formal semantics must give way to the quirks of individual psychology when discourse is explicitly addressed to that topic". From a psychological point of

view it is not a logical analysis of propositional attitudes that is called for, but one which refers to the actual way in which we handle propositional attitudes of others and of ourselves. Essential to Johnson-Laird's approach is an extension of the notion of a mental model that also incorporates what is known about the knowledge and beliefs of other speech participants. For that influences both production and interpretation. For example, how a speaker intends a description to designate and how she intends it to be taken by the hearer determines the way in which a certain belief may be reported. The general principle that seems to be at work here is that one should not use ways of designation that might implicate that one has knowledge one does not have. (This proposal has a distinct Gricean flavor, though Johnson-Laird does not make the connection.) Concerning substitution of equivalents, Johnson-Laird remarks that such inferences are warranted only if the model one has of someone's information enables one to actually draw the inference in question. Important in this respect is that the representation one has of someone else's beliefs may also contain information about misconceptions and the like, e.g. ones that concern his knowledge of the language. *My friend believes that transvestites are monks*, may report either a misconception of what transvestites are, i.e. of why some people behave in certain ways, or a misconception of what *transvestite* means. According to Johnson-Laird such examples show that a semantic theory in the logical vein, can never be rich enough to cope with propositional attitudes. People often have an imperfect grasp of their language and a good theory should be able to incorporate this phenomenon.

Johnson-Laird's conclusions are rather optimistic. If one is willing to barter the strongly anti-psychologistic realism of old for a more constructive approach, which is weakly mentalistic, it seems that the main obstacles to a psychologically plausible interpretation of the ways and means of formal semantics can be removed.

Max Cresswell's contribution is entitled 'The Autonomy of Semantics'. His aim is to investigate whether an autonomous discipline of semantics is possible. As for semantics, Cresswell regards adherence to the following principle as a *conditio sine qua non*. It is what he calls "the most certain principle" about meaning and it says that if of two sentences one is false and the other is true, then the two differ in meaning. This makes truth the central notion of semantics. Possible worlds semantics is a natural, though not the only possible development of this principle. The problem of propositional attitudes set aside, possible worlds semantics is indeed an autonomous discip-

line which relates to psychology as a special science in much the same way as psychology relates to physics. Though every psychological process and state involves physiological processes and states, this does not mean that the laws and concepts of psychology need to be reducible to those of physiology or physics. Analogously, Cresswell argues, though meanings are undoubtedly represented in the mind, truth-conditional semantics is autonomous in this sense that it does not involve this 'language of thought' in an essential way. This nice picture is disturbed, however, by the problem of propositional attitudes: (logically) equivalent propositions are interpreted as the same set of worlds (the same proposition) but people typically fail to behave accordingly. In the remainder of his paper Cresswell briefly reviews various solutions which have been proposed to this problem, and sketches one that would allow an autonomous semantics based on possible worlds, and which keeps as closely as possible to the ideal situation in which all language users are logically perfect.

The idea of taking propositions to be primitive entities, Cresswell rejects because he feels it contravenes his most certain principle: no essential relation between propositions and truth is involved. The problem with the impossible worlds approach is that it really comes down to using reinterpreted, weaker notions of e.g. necessity, than the usual ones. Quotational theories are rejected because they rest on an unanalyzed notion of meaning.

Another theory is Stalnaker's, who holds that in certain cases the objects of belief are propositions about the meanings of expressions. This idea reappears in Cresswell's sketch of his own solution. He observes that objects of belief in some sense have to be public entities since we refer to them in reporting belief. Cresswell argues for letting structured meanings, obtained from a sentence by replacing its symbols by their intensions, play this part. They are public, since they are derived from public objects (sentences), they do not contain anything we do not need anyway, and they fit into the truth-conditional framework. The problem of propositional attitudes can be solved, so it seems, by letting *believe* operate, not on the intension of its argument, but on the structured meaning. There is one problem that needs to be solved that has to do with iteration as in *Stephen believes that Veronica believes that Yvonne sings*. In order to prevent self-application of the intension of *believe*, which is set-theoretically impossible, Cresswell proposes to distinguish different 'levels of construal'. These may be needed also for non-iterated constructions in order to assign two different readings to e.g. *Belinda believes that the set of stars is finite*, one in which *believe* is an ordinary intensional operator, and hence closed under logical equivalence, and one in which it is a hyperinten-

sional operator which operates on a structured meaning. This leads to the postulation of a hierarchy of *believe*-operators with different, but related meanings which all depend on the meaning of the ordinary intensional *believe*. Cresswell notes that other solutions, too, implicitly use, or need, some kind of hierarchy.

The conclusions Barbara Partee reaches in her paper 'Belief Sentences and The Limits of Semantics', differ from Cresswell's. She claims that formal semantics ans psychological semantics are compatible once we give up the assumption that a human language must have a fixed and finitely representable semantics. That we should do so, she concludes from a discussion of what semantic competence should and could be, and of the semantics of belief-sentences.

What makes the problem of belief-sentences so fundamental for the formal semantics enterprise known as 'Montague Grammar' is that it seems to discredit it as a theory of the knowledge people have of their language. That substitution of logical equivalents fails to be truth-preserving in belief-contexts, Partee argues, means that we do not fully know our grammar. For equivalent sentences express the same proposition, but equivalence can be assumed to be undecidable and hence no finite being will be able to recognize every pair of equivalent sentences as such. This also brings out the relationship between semantic competence and the semantics of belief-sentences: it is our finiteness which both prevents us from knowing the semantics of our language completely and which causes substitution to fail in propositional attitude contexts. Regarding Montague Grammar as a theory of 'supercompetence' does not solve the problem, for even a supercompetent speaker will not substitute equivalents within a belief-ascription to other people. In an earlier paper Partee suggested that a more psychologistic approach to meanings might be helpful, and in order to gain some clarification she investigates what the adequacy criteria of various approaches to meaning imply a semantics of belief-sentences should look like.

In linguistics, synonymy judgments form the main database. But in this case they seem to provide only negative evidence: neither truth-values, nor truth-conditions discriminate finely enough. The adequacy criteria of procedural semantics, which sets out to model actual psychological processes, do lead to some positive requirements, but these concern the form of the theory (here Partee refers to the paper by Johnson-Laird). In formal semantics, entailment is the touchstone for theories, but in the case of belief-sentences positive evidence seems to be lacking. Though certain 'simple' entailments seem to be acceptable, the transitivity of entailment immediately produces

unacceptable results as well. Valid entailments abound, however, once auxiliary premisses are added. If some equivalence is believed then substitution is of course unobjectionable. Together with the procedural requirements these entailments form the positive material that a semantic theory of belief-sentences should account for.

A big problem in trying to reconcile psychological and formal semantics is to find a notion of semantic competence that will serve both. In formal semantics competence turns around entailment. The undecidability of the entailment-relation, Partee argues, need not be a problem. Semantic competence may be likened to knowing a finite set of axioms and rules which uniquely determines an undecidable set of valid entailments. But it does show that we should not think of semantic competence as the ability to make entailment judgments. That is something which finite beings such as humans will not generally be able to do. It follows that the notion of semantic competence that fits formal semantics must be free of any essential reference to psychological processes such as understanding and judging. The question then arises as to whether such a notion could also fit a psychological theory. It will not, if one claims that our knowledge of meaning consists in the procedures actually used in production and interpretation. So it will not fit procedural semantics. But psychological semantics could also be taken to include a theory of our knowledge of meaning, and then the notion of semantic competence outlined is compatible both with psychological and with formal semantics.

However, Partee sees serious problems ahead for this notion of semantic competence since it presupposes a fixed and finitely statable semantics. This might not be possible at all. The problem is that certain 'theory-loaded' terms do not allow a sharp distinction between meaning and (individual) beliefs. The following two sentences illustrate this: *Thomason believes that semantics is a branch of mathematics*; *Loar believes that semantics is a branch of psychology*. These two sentences, Partee argues, do not report a difference of opinion about what *semantics* means. Rather they report a debate about what semantics is, about how to go about doing it, about "how to carve up our conceptual space". If one assumes some fixed interpretation of *semantics*, one is doomed to get at least one of these two sentences wrong, i.e. one will unavoidably ascribe to Thomason, or to Loar (or to both) a belief he does not have. The proper way to understand the sentences, Partee says, borrowing a term from Johnson-Laird's, is to construct a mental model of (part of) the mental model of Thomason, c.q. Loar. A similar phenomenon occurs in language acquisition. Learning a word is not always simply establishing a link

between the word and some already available entity. Often the entity needs to be constructed, and this construction may go on after the word as such has been acquired.

Partee's conclusion is that for languages containing 'theory-loaded' terms no fixed and finitely statable semantics is possible. This fact reflects our awareness of our finiteness and of the necessary incompleteness of our knowledge. As such this does not discredit the way in which formal semanticists actually work when they develop fixed semantic theories for fragments of natural language, but it does show that a grand unification of these into one fixed semantic theory will not be the result of this. The assumption that this is possible should be abandoned and abandoning it also shows that no real conflict between the goals of psychological and formal semantics exists.

In their paper 'Computable Models of Belief and the Semantics of Belief Sentences' Robert Moore and Gary Hendrix argue that the study of computational models of belief will shed light on the semantics of belief-sentences. Their starting point is that computational models, as they are being studied in cognitive science, may serve to clarify conceptual problems. By developing a computational model that satisfies certain pre-theoretic notions, one may study their implications by investigating the consequences in terms of the model. This use of computational models they want to distinguish explicitly from computational *theories*, by which they understand theories about (cognitive) processes which claim that these processes are computational. The use of computational models does not involve that assumption.

A basic assumption of their model of belief is that belief is to be explained in terms of expressions of some kind of internal language. Immediately one thinks of philosophical arguments against the possibility of private languages. Moore & Hendrix discard these arguments by pointing to the internal languages of computer systems. The arguments brought forward should apply here as well, they claim, and thus are invalidated empirically.

In the computational model Moore & Hendrix sketch, belief consists in being in a computational relation to expressions of an internal language. These expressions form what they call a 'belief set'. The expressions in the belief set are assumed to be stored explicitly in memory, of which they form one space, other spaces corresponding to other attitudes. For the internal language Moore & Hendrix assume the language of predicate logic with propositional attitude operators. Also there are inference procedures which are at least capable of generating the usual valid inferences.

This simple computational model already explains various things Moore

& Hendrix claim, for example why there is no problem about logical consequence. If beliefs are individuated as formulas of the internal language some formula P might be contained in the belief set without necessarily some equivalent formula Q being in there as well. Although P and Q are equivalent according to the inference rules, the system may simply not have tried to derive Q from P, or its heuristics for applying the inference may fail to find a derivation, or the derivation may be too long to carry out.

So Moore & Hendrix conclude, the computational model clarifies some interesting conceptual puzzles. Also it will give an adequate semantics of belief-sentences. A first, rough formulation of this semantics is the following: 'A believes that S' is true iff the representation of S in the internal language is an element of A's belief set (or can be deduced from it with limited effort). What remains to be done, Moore & Hendrix note, is to specify the relation between S and the internal language expression representing S. Simply stating that there is an internal expression that A would express as S will not do. For one thing, there is *de re* belief. Also, this would imply that one will always be able to express one's beliefs in the language in which they are attributed. Ascription of belief, as Moore & Hendrix want to define it, must rest on some notion of sameness of meaning across individuals. They have an elaborate proposal for its definition, which uses, among others, the notion of a structured meaning à la Lewis for complex expressions. In terms of this definition they define under what circumstances a natural language expression can be said to express the same meaning as an internal expression for some individual. Then they state the truth conditions for *de dicto* and *de re* belief. 'A believes that S' is true under a *de dicto* interpretation iff A has some internal language expression P in his belief set such that S expresses the meaning of P for A. *De re* belief ascriptions Moore & Hendrix propose to handle as follows. Whereas a *de dicto* belief ascription uses a sentence which gives the meaning of some internal language expression in the belief set of the subject, a *de re* belief ascription uses a sentence of which some part gives, not the meaning of the corresponding part of the internal language espression, but its reference, with the added condition that the subject be able to identify this referent.

Moore & Hendrix conclude with some speculations about why so far no one has formulated a semantics of belief sentences in terms of psychological states and processes. For this they see two reasons. One is the idea that semantics should be autonomous, and should give the truth-conditions of sentences without recourse to how and why we understand and believe them. Moore & Hendrix object to this view by pointing out that psychological

states, such as belief, are what propositional attitude verbs, such as *believe*, *refer* to. Another motivation is the desire to equate knowledge of truth conditions with semantic competence. This, they feel, is mistaken: the actually true theory of some domain determines the truth conditions of sentences about it, and this theory is generally not known. Semantic competence consists in knowing which formula of the internal language represents a sentence of the natural language.

According to Janet Dean Fodor psychologists and linguists have a common goal, namely the development of a model of the knowledge individuals have of their language and of the way in which they use it. She assumes that it will be possible to develop one system of semantic representation that will serve both, and her paper 'The Mental Representation of Quantifiers' is meant as a contribution to this. A system of semantic representation should do various things, such as getting all the interpretations and entailments right, but Fodor's main concern is that it predict that there are marked differences between such interpretations and entailments. Some are harder to get than others. For example a sentence like *a child saw a squirrel*, which contains two occurrences of the same quantifier, is easier to interpret than *a child saw every squirrel*, which contains two occurrences of different quantifiers, the latter being ambiguous between two readings: the $\exists\forall$-reading and the $\forall\exists$-reading. Of these two, the latter is marked, and harder to get than the former. Here, clearly both the syntactic form and the form of the semantic representation play a role.

Using these, and other phenomena Fodor reviews some existing linguistic systems of semantic representation, and concludes that they fail to account for all of them. The systems she discusses are 'prefixed quantifier systems', which use standard logic at a level of logical form, 'hierarchical systems', and 'feature systems'.

Fodor's own system of semantic representation rests on the observation that quantifiers fall into two classes: those such as *all*, *many*, *seventeen* which are multiply instantiated, and those such as *a*, *some*, *the* which get a singular instantation. This is reflected in the way in which these quantifiers are represented in the diagrams which Fodor calls 'models-of-the world representations'. *Every child saw John* will be represented as in (a), and *a child saw John* as in (b):

(a)

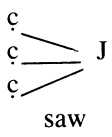

(b)

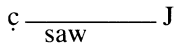

In order to distinguish between different multiply instantiated quantifiers (e.g. *every* and *three*) such models need to be supplemented with 'mental foototes' specifying number (or proportion). These models-of-the world representations, Fodor emphasizes, should not be confused with the logician's notion of a model. They are expressions of a certain formalism, that themselves need to be interpreted. Part of the adequacy of this system lies in the fact that this interpretation can be done in a psychologically more plausible way.

First of all, all phrases that are multiply instantiated, either because they contain a quantifier that requires this, or because they occur in the scope of such a quantifier, are represented alike, and differently from those that get a singular instantation. Another advantage is that quantified phrases are interpreted *in situ*. Sentences are interpreted 'on line', so no movement is assumed to take place. This explains e.g. why some readings are more difficult than others. The unmarked (∃∀-)reading of *a child saw every squirrel* is represented by first giving *a child* a singular instantation, and then *every squirrel* multiple instantation, as in diagram (c):

(c)

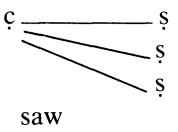

In order to get the marked ∀∃-reading, the representation of *a child* needs to be revised, since being in the scope of *every squirrel* a multiple instantiation is forced upon it. This means that the representation begun under (d)(i) should be continued, not as in (c), but as in (d)(ii):

(d)(i)

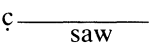

(d)(ii)

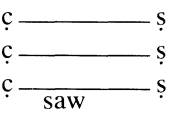

This system of representation, Fodor claims, is useful for other purposes as well. For example, collective readings are obtained if multiple paths converge before the representation of the verb, as in (e), which represents the collective reading of *all the children lifted the rock*:

(e)

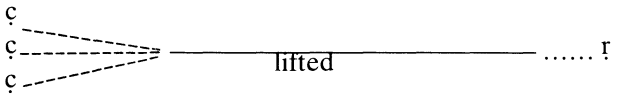

This 'stretching' of the notion of scope so as to include scope over a verb or other types of non-quantified constituents also explains well-known distributive differences between *each*, *every* and *all*. Discussing various examples, Fodor concludes that the rather squishy notion of a 'sphere of influence' of a quantifier in a certain position gives a better explanation of the facts than strict constraints on rules, such as the island-constraints.

She also acknowledges the existence of an absolute constraint on the relation between syntactic structure and semantic representation, which is that multiplicity of interpretation can only be inherited through a direct link with a multiple quantifier.

This constraint forces a particular representation for sentences containing three-place verbs. All arguments should be accessible to each other without intervention from the other, since e.g. a multiple subject term may bestow a multiple interpretation on the indirect object even if the direct object is singularly instantiated. Such representations then take the following form (f) and multiplicity of instantiation is then represented two-dimensionally, as in (g):

(f)

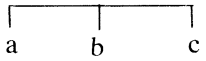

(g)

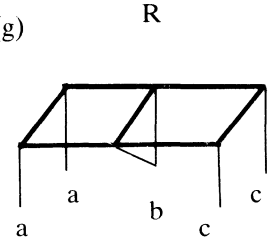

Fodor's main conclusion is that whereas linguists have not been concerned with the form in which information is represented because what they took these representations to look like did not tie up with their own syntactic representations, her own system of representation is not only plausible psychologically, but also close enough to linguistic representations so as to make linguists, psycholinguists and cognitive psychologists meet half-way.

Nuel D. Belnap Jr.'s paper 'Questions and Answers in Montague Grammar' is mainly concerned with a detailed analysis of several phenomena concerning the semantics of questions and the relationship between questions and answers. In an introductory section Belnap gives a brief, but forceful statement of how he sees the relation between formal theories of language, such as Montague Grammar, on the one hand and theories about natural language processing on the other. The essence of Montague Grammar, Belnap says, is that it gives a completely rigorous treatment of the syntax and seman-

tics of (some fragment of) some natural language. Such formal theories may be tied up with natural language processing in two ways. First of all, one might interpret the formal theory as the one which is actually used by people (or computers) as they process natural language. But there is also a second interpretation possible, which Belnap clearly favours. That is that we regard the formal theory as a procedure-independent description of the tasks a theory of processing should perform. Viewed in this way, the formal theory should preferably *not* refer to processing, psychological reality, and the like. The only constraint that a formal theory should meet is that it be a correct and rigorously formal description.

One of the objections one has heard raised against the application of model-theoretic semantics to natural language is that such a semantic theory is restricted by its very essence to just one type of language: descriptive language. In view of such objections, providing a sound model-theoretic semanics for questions in a natural language is an indirect, but important contribution to the discussion that is the main theme of this volume. It shows that there is no need to reject model-theoretic semantics from a psychological of linguistic point of view because its scope is limited.

Beside some powerful polemic remarks, Belnap's paper, which is based on joint work by himself and the late Michael Bennet, contains a detailed analysis of certain constructions involving questions, and many interesting observations, most of them concerning the relationship between questions and answers. Belnap argues, quite extensively, that an interesting systematic theory of the syntax and the semantics of answers is possible, if we are willing to distinguish between what he calls 'answers' and what are merely 'helpful responses'. This distinction, Belnap claims, can be made quite systematically. To a question such as *Which person kicked Sam?* a reaction like *I don't know*, or *Ask Sam*, is merely a helpful response.

One of the most interesting phenomena to which Belnap draws attention in this paper is the existence of such questions as *Where do two unicorns live?* This question, Belnap argues, has two different readings. One asks to give the singular place where two unicorns live. The second asks to give for two unicorns the place where each lives.

The semantics of such questions also has implications for what may count as an answer to them. On the wide scope reading of *two unicorns*, the question has, Belnap claims, more than one complete and true answer. Each true specification of the dwelling-places of some two unicorns is a complete and true answer, and of course there can be lots of those. Various frameworks and

analyses are criticized by Belnap for not allowing for this phenomenon. Such theories commit what he calls 'The Unique Answer Fallacy'.

A large part of Belnap's paper is devoted to a statement of explicit syntactic and semantic rules that will permit the derivation of questions on such readings. The process defined by Belnap resembles the standard process of quantification that a Montague Grammar uses to handle scope ambiguities, but the formal details are much more complicated and the reader is referred for their precise statement to the paper itself. There is also a detailed account of how to derive multiple questions, such as *Which woman loves which man?*

The paper ends with a section that contains some interesting observations which are intended to show that the semantics of various *wh*-words (*who*, *which person*, *what*, *which thing*) is more complicated than is sometimes assumed. Belnap argues here for example that *what* cannot simply be paraphrased as *which things*. For *What is in the basket?* and *Which things are in the basket?* differ at least in sofar as the former, but not the latter can be fully answered by *Three apples*.

The last paper in this volume, 'Linearization Describing Spatial Networks' by Willem J. Levelt, is a case study on a psycholinguistic subject, viz. the way in which speakers linearize their information. The paper does not address the main topic of the volume explicitly.

A fundamental characteristic of spoken language is its temporal, linear, left-to-right ordering. If information is to be encoded in speech this means that a choice has to be made about the ordering. This is neither a trivial nor a random matter. Sometimes the information itself is already linearly ordered, as for example in the report of an accident. In speech this ordering may then be taken over. But often information is not intrinsically ordered, e.g. when one gives a description of a room, and then a choice has to be made.

Linearization, in Levelt's view, is a process that is not fully determined by its content. It has certain independent functional properties. Levelt's conjecture is that general principles, also valid in other domains of discourse, underlie the way in which information is ordered in descriptions. The most general principle Levelt considers is the 'principle of minimal effort': everything else being equal, speakers will choose descriptions which minimize the number and duration of elements in store in memory, and the length of the description.

This principle is made more specific in the experiment described in Levelt's paper. People are asked to describe simple spatial networks: linear ones, hierarchical ones, and networks containing loops. Two models are

devised, one a speaker-orientated model and another a listener-orientated model. They are given in the form of an ATN. The main difference between the two is that in the first one, the description 'jumps' back to an unfinished choice-point (a point in the network from which more than one branch departs), whereas in the second it 'moves' back by explicitly retracing its steps.

From the principle of minimal effort three global constraints are derived. One is that short branches are preferred to long ones, the second that embedding of choice-points is minimized, and the third that loops are described before other branches. The experimental findings which Levelt reports strongly support the existence of these constraints. Moreover the two types of linearization embodied in the two models are indeed the two types to be found among actual speakers.

Levelt classifies linearization as a non-linguistic process. Whether it interacts in any way with for example the relations between syntactic structure and semantic structure needs to be further investigated.

As the review of the individual papers may have made clear, they contain diverging views on the main theme of this volume, namely the relation between psychological and formal semantics. Some positions are argued for in more detail than others, the amount of detail perhaps depending on the extent to which the position is an accepted one. The papers of Partee and especially Johnson-Laird, who argue for a closer cooperation, contain detailed argumentations for what from both perspectives is a new enterprise. The thesis that semantics is an autonomous discipline, on the other hand, defended by Cresswell and Belnap, is rather common among semanticists of a philosophical and logical inclination. The defence Cresswell offers in his paper remains rather sketchy. The account of the semantics of propositional attitude verbs that it presupposes is outlined, but not worked out in any detail, and so the strength of the defence is hard to judge on the basis of his paper. The reader must find his or her way to the details of the theory through the references given in Cresswell's paper. As for Belnap's argumentation, it would be rather interesting to explore one of his (implicit) arguments for autonomy that relates to the semantics of questions and answers, viz. that a systematic difference can be made between answers, which can be characterized semantically, and responses, which belong to the domain of what Belnap calls 'conversation theory'. It seems that the resulting theory of answerhood is rather limited. Using a more liberal notion, at the risk of making semantics less autonomous, might yield a more encompassing theory.

Moore & Hendrix, who argue for a psychological theory of meaning, are explicit about their assumptions and their analysis is rather detailed, but it remains unclear whether their semantical analysis of belief-sentences meets the requirements of the formal semanticists. Their theory resembles a syntactic approach and the well-known problems of such an approach are various. Moore & Hendrix discuss the problem of logical equivalence noting that if beliefs are individuated as formulas, and a belief set is taken to be a set of formulas, the equivalence of P and Q need not imply that if one of them is in a belief set, the other is too. On the other hand, they also state that for a belief ascription to be true, it is required that a certain formula is in an individual's belief set or 'can be derived in his belief set with limited effort'. But this is precisely what has worried formal semanticists: how to account for the fact that people are not logically omniscient, but are not completely logically ignorant either. Moore & Hendrix do not make clear that their system is one that allows one to deal with this.

Analogous remarks can be made about Fodor's semantic representation system for quantifiers. One of the requirements such systems should meet, she says, is that it will enable one to represent all interpretations, and account for all entailments, which is typically one of the main concerns of a formal semanticist. Whether Fodor's system fulfils this requirement remains unclear. As a matter of fact, whether it is a system of semantic representation as a formal semanticist would define it, rather than one of syntactic paraphrase, is not clear either. The purported relation between Fodor's models-of-the-world and the 'real' world is not defined. And its definition will not be a straightforward application of known techniques. For one thing, the mixture of two representational devices that Fodor make use of is not something one finds in ordinary model-theoretic semantics.

A last remark that perhaps needs to be made concerns the interpretation that is given of the term 'formal semantics'. Often, for example in Johnson-Laird's paper, it is characterized in such a way that it coincides with intensional, possible worlds semantics. This surely is the best known and most widely used theory and there is nothing wrong with concentrating on it. But it should be borne in mind that intensional semantics is not the only variety of formal semantics that is possible. In fact, various other approaches are currently being developed which are intended to overcome some of the difficulties, such as completeness, that stand in the way of a psychologically plausible interpretation of intensional semantics.

However, none of these remarks should lead the reader to believe that the papers in his volume are not very worthwhile reading for anyone who is interested in the semantics of natural language. Clearly not all is said and done. Many questions remain unanswered and many problems unsolved. But new questions are raised and new light is shed on old ones. And that is the hallmark of a valuable work.