



The Structure of a Semantic Theory

Jerrold J. Katz; Jerry A. Fodor

Language, Vol. 39, No. 2. (Apr. - Jun., 1963), pp. 170-210.

Stable URL:

<http://links.jstor.org/sici?sici=0097-8507%28196304%2F06%2939%3A2%3C170%3ATS0AST%3E2.0.CO%3B2-J>

Language is currently published by Linguistic Society of America.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/lsa.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact support@jstor.org.

THE STRUCTURE OF A SEMANTIC THEORY

JERROLD J. KATZ

JERRY A. FODOR

Massachusetts Institute of Technology

1. Introduction. This paper¹ does not attempt to present a semantic theory of a natural language, but rather to characterize the form of such a theory. A semantic theory of a natural language is part of a linguistic description of that language. Our problem, on the other hand, is part of the general theory of language, fully on a par with the problem of characterizing the structure of grammars of natural languages. A characterization of the abstract form of a semantic theory is given by a metatheory which answers such questions as these: What is the domain of a semantic theory? What are the descriptive and explanatory goals of a semantic theory? What mechanisms are employed in pursuit of these goals? What are the empirical and methodological constraints upon a semantic theory?

The present paper approaches the problem of characterizing the form of semantic theories by describing the structure of a semantic theory of English. There can be little doubt but that the results achieved will apply directly to semantic theories of languages closely related to English. The question of their applicability to semantic theories of more distant languages will be left for subsequent investigations to explore. Nevertheless, the present investigation will provide results that can be applied to semantic theories of languages unrelated to English and suggestions about how to proceed with the construction of such theories.

We may put our problem this way: What form should a semantic theory of a natural language take to accommodate in the most revealing way the facts about the semantic structure of that language supplied by descriptive research? This question is of primary importance at the present stage of the development of semantics because semantics suffers not from a dearth of facts about meanings and meaning relations in natural languages, but rather from the lack of an adequate theory to organize, systematize, and generalize these facts. Facts about the semantics of natural languages have been contributed in abundance by many diverse fields, including philosophy, linguistics, philology, and psychology. Indeed, a compendium of such facts is readily available in any good dictionary. But at present the superabundance of facts obscures a clear view of their interrelations, while such theories as have been proposed to account for the facts have, in general, been either too loosely formulated or too weak in explanatory and descriptive power to succeed.

2. The projection problem. A full synchronic description of a natural language is a grammatical and semantic characterization of that language (where the term 'grammatical' is construed broadly to include phonology, phonemics,

¹ This work was supported in part by the U. S. Army Signal Corps, the Air Force Office of Scientific Research, and the Office of Naval Research; and in part by the National Science Foundation (Grant G-13903).

morphology, and syntax). Hence, a semantic theory must be constructed to have whatever properties are demanded by its role in linguistic description. Since, however, the goals of such description are reasonably well understood and since, in comparison to semantics, the nature of grammar has been clearly articulated, we may expect that by studying the contribution that semantics will be required to make to a synchronic description of a language we can clarify the subject, the form of generalizations, the goals, and the empirical and methodological constraints upon a semantic theory.

A fluent speaker's mastery of his language exhibits itself in his ability to produce and understand the sentences of his language, INCLUDING INDEFINITELY MANY THAT ARE WHOLLY NOVEL TO HIM (i.e. his ability to produce and understand ANY sentence of his language²). The emphasis upon novel sentences is important. The most characteristic feature of language is its ability to make available an infinity of sentences from which the speaker can select appropriate and novel ones to use as the need arises. That is to say, what qualifies one as a fluent speaker is not the ability to imitate previously heard sentences but rather the ability to produce and understand sentences never before encountered. The striking fact about the use of language is the absence of repetition: almost every sentence uttered is uttered for the first time. This can be substantiated by checking texts for the number of times a sentence is repeated. It is exceedingly unlikely that even a single repetition of a sentence of reasonable length will be encountered.

A synchronic description of a natural language seeks to determine what a fluent speaker knows about the structure of his language that enables him to use and understand its sentences. Since a fluent speaker is able to use and understand any sentence drawn from the INFINITE set of sentences of his language, and since, at any time, he has only encountered a FINITE set of sentences, it follows that the speaker's knowledge of his language takes the form of rules which project the finite set of sentences he has fortuitously encountered to the infinite set of sentences of the language. A description of the language which adequately represents the speaker's linguistic knowledge must, accordingly, state these rules. The problem of formulating these rules we shall refer to as the projection problem.

This problem requires for its solution rules which project the infinite set of sentences in a way which mirrors the way that speakers understand novel sentences. In encountering a novel sentence the speaker is not encountering novel elements but only a novel combination of familiar elements. Since the set of sentences is infinite and each sentence is a different concatenation of morphemes, the fact that a speaker can understand any sentence must mean that the way he understands sentences which he has never previously encountered is compositional: on the basis of his knowledge of the grammatical properties and the meanings of the morphemes of the language, the rules which the speaker knows enable

² There are exceptions, such as sentences with technical words that the speaker does not know and sentences too long for the speaker to scan in his lifetime. But these exceptions are of no systematic importance. Analogously, a person's mastery of an algorithm for propositional calculus can be said to exhibit itself in his ability to decide mechanically whether ANY well-formed formula of propositional calculus is a tautology, even though some well-formed formulae are too long for human processing, etc.

him to determine the meaning of a novel sentence in terms of the manner in which the parts of the sentence are composed to form the whole. Correspondingly, we can expect that a system of rules which solves the projection problem must reflect the compositional character of the speaker's linguistic knowledge.³

3. Synchronic linguistic description minus grammar equals semantics. A description of a natural language is, *inter alia*, a solution to the projection problem for that language. If we are to discover the goals of semantics by subtracting from the goals of a description of a language whatever the grammar contributes to the solution of the projection problem, we must consider the respect in which a grammar is a solution for the grammatical aspect of the projection problem.

Grammars answer the question: What does the speaker know about the phonological and syntactic structure of his language that enables him to use and understand any of its sentences, including those he has not previously heard? They do so by providing rules which generate the sentences of the speaker's language. In particular, these rules generate infinitely many strings of morphemes which, though they are sentences of the language, have never been uttered by speakers.⁴ Moreover, a grammar generates the sentences which a speaker is, in principle, capable of understanding in such a way that their derivations provide their structural descriptions. Such descriptions specify the elements out of which a sentence is constructed, the grammatical relations between these elements and between the higher constituents of the sentence, the relations between the sentence and other sentences of the language, and the ways the sentence is syntactically ambiguous together with an explanation of why it is ambiguous in these ways. Since it is this information about a novel sentence which the speaker knows and which enables him to understand its syntactic structure if and when he encounters the sentence, an adequate transformational grammar of a language PARTIALLY solves the projection problem for the language.

A semantic theory of a language completes the solution of the projection problem for the language. Thus, semantics takes over the explanation of the

³ A solution to the projection problem is certainly less than a full theory of speech. In particular, it does not provide a theory of speech production (or recognition). The difference between a description of a language and a theory of speech production is the difference between asking for a characterization of the rules of language which a speaker knows and asking for an account of how he actually applies those rules in speaking. Some things left out by the first theory but not by the second are considerations of the psychological parameters of speech production (e.g. limitations of immediate memory, level of motivation), and developmental accounts of the way the child becomes a fluent speaker (by conditioning? by the exploitation of innate mechanisms? by some combination of innate endowment and learning?). Though such problems concerning speech production lie outside the scope of a theory of a language, such a theory is essential to a theory of speech production. It is first necessary to know *WHAT* is acquired and used before it is sensible to ask *HOW* it is acquired and used.

⁴ This conception of grammar is due to Chomsky. Cf. *Syntactic structures*² ('s-Gravenhage, 1962); 'Three models for the description of language', *I.R.E. transactions on information theory*: Vol. IT-2, Proceedings of the Symposium on Information Theory (Sept. 1956). For a bibliography on transformational grammar, see Chomsky, 'On the notion "rule of grammar"', *Structure of language and its mathematical aspects*: Proceedings of symposia in applied mathematics 12.16 fn. 24 (1961).

speaker's ability to produce and understand new sentences at the point where grammar leaves off. Since we wish to determine, when we have subtracted the problems in the description of a language properly belonging to grammar, what problems belong to semantics, we must begin by gaining some grasp of how much of the projection problem is left unsolved by an optimal grammar.

One way to appreciate how much of understanding sentences is left unexplained by grammar is to compare the grammatical characterizations of sentences to what we know about their semantic characterizations. If we do this, we notice that the grammar provides identical structural descriptions for sentences that are different in meaning and different structural descriptions for sentences that are identical in meaning. The former will be the case for all morphemically distinct substitution instances of a given sentential type; for example, *The dog bit the man* and *The cat bit the woman*. The latter will be the case for many instances of sentential synonymy; for example, *The dog bit the man* and *The man was bitten by the dog*.⁵

In general, it is obvious that in no sense of meaning does the structural description which the grammar assigns to a sentence specify either the meaning of the sentence or the meaning of its parts. Such considerations must now be made precise in order that we may apply our formula 'linguistic description minus grammar equals semantics' to determine a lower bound on the domain of a semantic theory. Later in this section we will fix an upper bound by determining what problems lie outside the concerns of a complete linguistic description.

Grammars seek to describe the structure of a sentence IN ISOLATION FROM ITS POSSIBLE SETTINGS IN LINGUISTIC DISCOURSE (WRITTEN OR VERBAL) OR IN NON-LINGUISTIC CONTEXTS (SOCIAL OR PHYSICAL). The justification which permits the grammarian to study sentences in abstraction from the settings in which they have occurred or might occur is simply that the fluent speaker is able to construct and recognize syntactically well-formed sentences without recourse to information about settings, and this ability is what a grammar undertakes to reconstruct. Every facet of the fluent speaker's linguistic ability which a grammar reconstructs can be exercised independently of information about settings: this is true not only of the ability to produce and recognize sentences but also of the ability to determine syntactic relations between sentence types, to implicitly analyze the syntactic structure of sentences, and to detect grammatical ambiguities. Since, then, the knowledge that a fluent speaker has of his language enables him to determine the grammatical structure of any sentence without reference to information about setting, grammar correspondingly forms an independent theory of this independent knowledge.

We may generalize to arrive at a sufficient condition for determining when an ability of speakers is the proper subject matter of a synchronic theory in linguistics. The generalization is this: IF SPEAKERS POSSESS AN ABILITY THAT ENABLES THEM TO APPREHEND THE STRUCTURE OF ANY SENTENCE IN THE INFINITE

⁵ Moreover, sentences that receive the same structural description may differ in that one is semantically ambiguous or anomalous but the other is not. Compare *The bill is large*, *The paint is silent*, and *The street is wide*, all of which receive the same structural description from the grammar.

SET OF SENTENCES OF A LANGUAGE WITHOUT REFERENCE TO INFORMATION ABOUT SETTINGS AND WITHOUT SIGNIFICANT VARIATION FROM SPEAKER TO SPEAKER, THEN THAT ABILITY IS PROPERLY THE SUBJECT MATTER OF A SYNCHRONIC THEORY IN LINGUISTICS.

The first question in determining the subject matter of a semantic theory is: Can we find an ability which satisfies the antecedent of this generalization, which is beyond the range of grammatical description, and which is semantic in some reasonable sense? If we can, then that ability falls within the domain of a semantic theory.

In order to find such an ability, let us consider a communication situation so constructed that no information about setting can contribute to a speaker's understanding of a sentence encountered in that situation. Any extragrammatical ability that a speaker can employ to understand the meaning of a sentence in such a situation will ipso facto be considered to require semantic explanation.

The type of communication situation we shall consider is the following. A number of English-speakers receive an anonymous letter containing only the English sentence S. We are interested in the difference between this type of situation and one in which the same anonymous letter is received by persons who do not speak English but are equipped with a completely adequate grammar of English. To investigate what the first group can do by way of comprehending the meaning of S that the second group cannot is to factor out the contribution of grammar to the understanding of sentences. We will only investigate aspects of linguistic ability which are invariant from individual to individual within each group. We thus make sure that the abilities under investigation are a function not of idiosyncrasies of a speaker's personal history but only of his knowledge of his language.

Suppose S is the sentence *The bill is large*. Speakers of English will agree that this sentence is ambiguous, i.e. that it has at least two readings. According to one it means that some document demanding a sum of money to discharge a debt exceeds in size most such documents; according to the other it means that the beak of a certain bird exceeds in bulk those of most similar birds. However, the fact that this sentence is ambiguous between these readings cannot be attributed to its syntactic structure, since, syntactically, its structure on both readings is as shown in Figure 1. That is, the group who do not speak English but are equipped with a grammar can say no more about *The bill is large* than what is represented in Fig. 1. Thus, this sentence, which is marked as unambiguous by




FIG. 1

the grammar, will be understood as ambiguous by a fluent speaker. From this difference between the performances of the two groups, it follows that one facet of the speaker's ability that a semantic theory will have to reconstruct is that he can detect nonsyntactic ambiguities and characterize the content of each reading of a sentence.

Now suppose S is the sentence *The bill is large but need not be paid*. Speakers of English will understand this sentence only on readings in which *bill* means an order to pay a sum of money to discharge a debt. This shows that a speaker can disambiguate parts of a sentence in terms of other parts and thereby determine the number of readings of a sentence. Thus, another facet of the speaker's semantic ability is that of determining the number of readings that a sentence has by exploiting semantic relations in the sentence to eliminate potential ambiguities.

Now let S be the sentence *The paint is silent*. English speakers will at once recognize that this sentence is anomalous in some way. For example, they will distinguish it from such sentences as *The paint is wet* and *The paint is yellow* by applying to it such epithets as 'odd', 'peculiar', 'paradoxical', and 'bizarre'. Though it is clear that the speaker does not have the explicit conceptual machinery to correctly characterize the difference between these sentences, his consistent use of such rough labels shows that he is aware of some sort of linguistic anomaly. But the group who do not speak English and are equipped only with a grammar will regard all these sentences as fully regular, since there is no grammatical basis for distinguishing between them. Hence, another facet of the semantic ability of the speaker is that of detecting semantic anomalies. Correspondingly, a semantic theory will be needed to mark the distinction between semantically anomalous sentences and semantically regular sentences, so far as this distinction is not co-extensive with the distinction the grammar makes between ungrammatical and grammatical strings of morphemes.

Finally, whatever sentence the anonymous letter contains, as a rule, speakers of English can easily decide what sentences are paraphrases of it and what are not, in the sense that they can answer the questions *What does the letter say? Does the letter say such-and-such? How can what the letter says be rephrased?* This facet of the speaker's ability cannot be referred to his mastery of grammar either, for a person who is equipped with a grammar but who does not speak English will be unable to tell whether or not a sentence is a paraphrase of S. The reasons are simply that there need be no definite grammatical relation between a sentence and its paraphrases, e.g. between *Two chairs are in the room* and *There are at least two things in the room and each is a chair*, and that where a definite grammatical relation obtains between a pair of sentences, neither need be a paraphrase of the other, e.g. *The ball was hit by the man* and *The ball was hit*, *The man hit the ball* and *The man did not hit the ball*.⁶ Thus, still another facet of the speaker's semantic ability which must fall within the domain of a semantic theory is his paraphrasing skill.

We can now tentatively characterize the lower bound on the domain of a

⁶ Cf. *Syntactic structures*, Appendix II, for the transformations which relate these sentences.

semantic theory, since we have found an ability of speakers which cannot be accounted for by grammar, which is semantic in a reasonable sense, and which enables speakers to apprehend the semantic structure of an infinite number of sentences without information about setting and independent of individual differences between speakers. We thus take the goals of a semantic theory to include at least the explication of each facet of this ability and of the interrelations between them.

The speaker's exercise of this ability, which henceforth we shall refer to as *THE ABILITY TO INTERPRET SENTENCES*, provides empirical data for the construction of a semantic theory, just as the construction of a grammar draws upon empirical data supplied by the exercise of the speaker's ability to distinguish well-formed sentences from ungrammatical strings, to recognize syntactic ambiguity, and to appreciate relations between sentence types. A semantic theory describes and explains the interpretative ability of speakers by accounting for their performance in determining the number and content of the readings of a sentence, by detecting semantic anomalies, by deciding on paraphrase relations between sentences, and by marking every other semantic property or relation that plays a role in this ability.

4. What is beyond the descriptive scope of a semantic theory. Having fixed a lower bound on the domain of a semantic theory, our next step must be to fix an upper bound, thus uniquely determining the set of problems forming the domain of a semantic theory of a natural language.

Previous conceptions of semantics have usually defined the goals of a semantic description of a natural language in such a way that to achieve them a semantic theory would have to account for the manner in which settings determine how an utterance is understood. We shall now show that to set the goals of a semantic theory this high is to set them too high. Once we have shown that a semantic theory cannot be expected to account for the way settings determine how an utterance is understood, we will have fixed an upper bound on the domain of semantic theories. That is, we will have shown that a semantic theory is a theory of the speaker's ability to interpret the sentences of his language.

The form of a theory of how settings control the understanding of utterances of sentences is as follows. Such a theory is a function F whose arguments are a sentence S ; GS , a grammatical description of S ; IS , a semantic interpretation of S (where IS is the set of possible readings of S); and C , an abstract characterization of a setting. $F(S, GS, IS, C)$ is

- (1) the particular reading in IS that speakers of the language give to S in settings of the type C , or
- (2) An n -tuple ($n \geq 2$) of the readings from IS that speakers of the language give to S if S is ambiguous n -ways in settings of type C , or
- (3) The null element if speakers of the language give to S none of the readings in IS when S occurs in settings of type C .

The value of $F(S, GS, IS, C)$ is (1) just in case C fully disambiguates S , i.e. C determines a unique reading from the one or more in IS ; it is (2) just in case C

fails to fully disambiguate S; it is (3) just in case an occurrence of S in C is token-odd.⁷

An example of each of these cases will clarify this abstract formulation by showing how a theory of this form would explicate the speaker's ability to choose the reading(s) which a setting determines for a sentence occurring in it. As an example of case (1) consider the sentence *The shooting of the hunters was terrible*. This sentence is ambiguous between the reading r_1 , on which it means that it was terrible that the hunters were shot, and the reading r_2 , on which it means that the marksmanship of the hunters was very bad. This ambiguity will be represented in IS. The theory F must decide which of these readings the sentence bears in settings which disambiguate it, and it must decide in which settings the sentence remains ambiguous. If, then, an utterance of the sentence occurs as an answer to the question *How good was the marksmanship of the hunters?*, i.e. if C represents a situation in which the marksmanship of the hunters is clearly at issue, then, *ceteris paribus*, the value of F would have to be r_2 .⁸ Now consider case (2). The ambiguous sentence *He follows Marx*, occurring in a setting in which it is clear that the speaker is discussing intellectual history, cannot bear the reading 'He dogs the footsteps of Groucho'. However, this setting leaves the sentence ambiguous between the readings 'He is a disciple of Karl's' and 'He postdates Karl'. Thus, F will have to have these latter two readings as its value for this sentence and this setting as arguments. Finally, as case (3), suppose the sentence *This is the happiest night of my life* is uttered in the middle of the day. Since the sentence is uttered in a setting that lacks conditions which utterances of this sentence presuppose, the occurrence is a case of token-oddity. Thus, for this sentence occurrence F must give the null element as its value, i.e. none of the readings of this sentence in IS are selected by C.

This, then, is the form of a theory about the effect of setting upon the way speakers understand sentences. Any particular theory is complete just to the extent that it solves the problems incorporated in this abstract formulation. A complete theory of this kind is more powerful in principle than a theory of the semantic interpretation of sentences in isolation. But a theory of settings must contain a theory of semantic interpretation as a proper part because the readings that a speaker attributes to a sentence in a setting are a selection from among those that the sentence has in isolation. It is clear that, in general, a sentence cannot have readings in a setting which it does not have in isolation. Of course, there are cases in which a sentence may have a reading for some speakers in some settings which it does not have in isolation for all speakers. But these cases are essentially idiomatic in the sense that meaning is determined either by special stipulation (passwords, nonce senses, etc.) or special rules (codes, etc.) or else by special information about the intentions of the speaker. If a theory of the selective effect of setting were required to deal with such cases, no such theory

⁷ Semantic type oddity is precluded by the assumption that IS contains at least one reading.

⁸ In the case where a sentence has exactly one reading in IS, i.e. is unambiguous, that reading must by the theory be assigned to the sentence in each and every normal setting.

would be possible, because any sentence may be made to mean anything you like simply by constructing the setting to include the appropriate stipulation.⁹ Since, then, the readings that a speaker gives a sentence in a setting are a selection from those which it has in isolation, a theory of semantic interpretation is logically prior to a theory of the selective effect of setting.

The abstract formulation given above may be realized in the form of a theory of either of two kinds, depending on how the notion of setting is construed. One kind of theory of setting selection construes the setting of an utterance to be the nonlinguistic context in which the utterance occurs, i.e. the full sociophysical environment of the utterance. The other kind takes the setting of an utterance to be the linguistic context in which the utterance occurs, i.e. the written or spoken discourse of which the utterance is a part. We shall consider, in turn, the possibility of constructing a theory of each of these types.

The first kind of theory of setting selection seeks to account for the way in which aspects of the sociophysical world control the understanding of sentences. Differing varieties of this kind of theory may be obtained by varying the aspects of the sociophysical environment of which the rules of the theory are permitted to take account, and by varying the spatiotemporal parameters of the environment. But clearly a necessary condition which any variety of this kind of theory must satisfy is that its construction of setting is so defined that it is able to represent all the nonlinguistic information required by speakers for understanding sentences. So far as a theory fails to satisfy this condition, it is incomplete, since there is then some information which determines the way speakers understand a sentence but which the theory fails to represent as part of the setting of that sentence.

But a complete theory of this kind is not possible in principle; for to satisfy the above necessary condition it would be required that the theory represent ALL the knowledge speakers have about the world. That this is so can be seen from even a few examples which show how nonlinguistic information of any kind may be involved in the understanding of a sentence. Consider (1) *Our store sells alligator shoes* and (2) *Our store sells horse shoes*. In normal settings (e.g. as signs in a store window or as newspaper advertisements), occurrences of (1) will be taken on the reading 'our store sells shoes made from alligator skins' while (2) will be taken on the reading 'our store sells shoes for horses'. Notice, however, that (1) is open to the reading 'our store sells shoes for alligators' and (2) is open to the reading 'our store sells shoes made from the skin of horses'. From this it follows that if a theory of setting selection is to choose the correct reading for (1), it must represent the fact that, to date, alligators do not wear shoes, although shoes for people are sometimes made from alligator skin. Conversely, if the theory is to choose the correct reading for (2), it must represent the fact that horses wear shoes, although shoes for people are not usually made from the

⁹ Take the following example. Let m be a one-to-one mapping of the set of English sentences onto itself such that the image of each sentence is a sentence which differs from it in meaning. Then the sentence *The sentence S which immediately follows this sentence is to be understood as $m(S)$* is a setting such that the meaning of a sentence occurring in it is not one of the meanings of that sentence in isolation.

skin of horses. Other examples illustrate the same point. Compare the three sentences: *Should we take junior back to the zoo? Should we take the lion back to the zoo? Should we take the bus back to the zoo?* Information which figures in the choice of the correct readings for these sentences includes the fact that lions, but not children and busses, are often kept in cages. Three further cases of the same sort are: *Can I put the wall-paper on? and Can I put the coat on?; Joe jumped higher than the Empire State Building and Joe jumped higher than you; Black cats are unlucky and People who break mirrors are unlucky.*¹⁰

The reader will find it an easy matter to construct an ambiguous sentence whose resolution requires the representation of practically any item of information about the world he chooses.¹¹ Since a complete theory of setting selection must represent as part of the setting of an utterance any and every feature of the world which speakers need in order to determine the preferred reading of that utterance, and since, as we have just seen, practically any item of information about the world is essential to some disambiguations, two conclusions follow. First, such a theory cannot in principle distinguish between the speaker's knowledge of his language and his knowledge of the world, because, according to such a theory, part of the characterization of a LINGUISTIC ability is a representation of virtually all knowledge about the world that speakers share. Second, since there is no serious possibility of systematizing all the knowledge of the world that speakers share, and since a theory of the kind we have been discussing requires such a systematization, it is ipso facto not a serious model for semantics. However, none of these considerations is intended to rule out the possibility that, by placing relatively strong limitations on the information about the world that a theory can represent in the characterization of a setting, a LIMITED theory of selection by sociophysical setting can be constructed. What these considerations do show is that a COMPLETE theory of this kind is impossible.

The second kind of realization of the abstract formulation of a theory of setting selection is one in which the setting of an occurrence of a sentence is construed as the written or spoken discourse of which the occurrence is a part. Such a theory has a strong and a weak version. The strong version requires that the theory interpret a discourse in the same way that a fluent speaker would (i.e. mark the ambiguities that the speaker marks, resolve the ambiguities that the speaker resolves, detect the anomalous strings that the speaker detects, recognize paraphrase relations that the speaker recognizes, and do all this both within and across sentence boundaries). Since, however, in so interpreting a discourse a speaker may need to bring to bear virtually any information about the world that he and other speakers share, the argument given against a complete theory of selection by sociophysical setting applies equally against the strong version of a theory of selection by discourse. Hence we need only consider the weak version.

¹⁰ We express our gratitude to David Bellugi for referring us to *My little golden book of jokes* (New York, 1961), from which these examples are drawn.

¹¹ We have convinced ourselves of the truth of this claim by making it the basis of a party game. One person supplies a fact, however obscure, and the others try to construct a sentence which that fact disambiguates. The game is not remarkably amusing, but it is surprisingly convincing.

The weak version of such a theory requires only that the theory interpret discourses just so far as the interpretation is determined by grammatical and semantic relations which obtain within and among the sentences of the discourse. Thus, such a theory seeks to disambiguate sentences and sequences of sentences in terms of grammatical and semantic relations between these and the sentences which form their setting in a discourse, to determine when an occurrence of a sentence or of a sequence of sentences is rendered anomalous by the sentences which form its setting in a discourse, and to recognize paraphrase relations between pairs of sentences and pairs of sequences of sentences in a discourse.¹²

But it is not at all clear that the weak version of a theory of discourse setting selection has greater explanatory power in these respects than a theory of semantic interpretation, since except for a few types of cases (see below), a discourse can be treated as a single sentence in isolation by regarding sentence boundaries as sentential connectives. As a matter of fact, this is the natural treatment. Consider the two-sentence discourse: *I shot the man with a gun. If the man had had a gun too, he would have shot me first.* The first sentence of this discourse is ambiguous in isolation, but not in this setting. But the problem of explaining this disambiguation is the same as the problem of explaining why the single sentence *I shot the man with a gun, but if the man had had a gun too, he would have shot me first* does not have an ambiguous first clause. This technique of replacing discourses or stretches in discourse by single compound sentences, using sentence connectives in place of sentence boundaries, clearly has a very extensive application in reducing problems of setting selection to problems of semantic interpretation of sentences in isolation. Thus, given a theory of semantic interpretation, little is left for a theory of setting selection to explain.

The fact which underlies this technique is that, in the great majority of cases, the sentence break in a discourse is simply equivalent to the conjunction *and*. (In others it is equivalent to *but*, in others to *for*, in others to *or*, etc.) Sometimes, however, a discourse cannot be directly converted into a compound sentence in this way. For example, the discourse *How are you feeling today? I am fine, thanks.* does not convert to **How are you feeling today and I am fine, thanks.* because the compound sentence is ungrammatical. But the fact that sentences of different types cannot be run together in the obvious way may not pose a serious problem; for it is not at all clear that less obvious conversions will not lead to a satisfactory treatment of such cases within a theory of semantic interpretation. For example, we may convert the discourse just cited into the single sentence *X asked 'How are you feeling today?' and Y replied 'I am fine, thanks.'* If such conversions can be carried out generally, then any problem about disambiguation, detection of anomaly, etc. that can be raised and/or solved in a theory of setting selection can be raised and/or solved by reference to an analogue in the theory of semantic interpretation. But even if such conversions cannot be carried out generally, the most interesting and central cases will still be within the range of a theory of semantic interpretation. Hence, for every discourse there is a single sentence

¹² For examples of studies toward a theory of this kind, cf. Z. S. Harris, 'Discourse analysis', *Lg.* 26.1-30 (1952); H. Herzberger, *Contextual analysis* (Princeton University dissertation 1957).

which consists of the sequence of n sentences that comprise the discourse connected by the appropriate sentence connectives and which exhibits the same semantic relations exhibited in the discourse. But since the single sentence is, *ex hypothesi*, described by a theory of semantic interpretation, in every case in which a discourse can be treated as a single sentence, a theory of semantic interpretation is descriptively as powerful as a theory of setting selection.

We opened the discussion of theories of setting selection in order to fix an upper bound on the domain of a semantic theory of a natural language. The result of the discussion is that, where such a theory is not reducible to a theory of semantic interpretation, it cannot be completed without systematizing all the knowledge about the world that speakers share and keeping this systematization up to date as speakers come to share more knowledge. A limited theory of how sociophysical setting determines the understanding of an utterance is possible, but even such a theory blurs the distinction between the speaker's knowledge of his language (his linguistic ability) and his knowledge of the world (his beliefs about matters of fact). Therefore, since it is unlikely that anything stronger than a theory of semantic interpretation is possible and since such a theory is clearly an essential part of a linguistic description, it is reasonable to fix the upper bound of a semantic theory of a natural language at the point where the requirements upon a theory of semantic interpretation are satisfied.

5. The components of a semantic theory. We must now determine what mechanisms a semantic theory employs in reconstructing the speaker's ability to interpret sentences. We have seen that this ability is systematic in that it enables the speaker to understand sentences he has never heard before and to produce novel sentences that other speakers understand in the way that he understands them. To account for this ability a semantic theory must be so formulated that its output matches the interpretive performance of a fluent speaker. In this section, we describe the form of semantic theories.

It is widely acknowledged and certainly true that one component of a semantic theory of a natural language is a dictionary of that language. The reason for including a dictionary as a component of a semantic theory is based on two limitations of a grammatical description. First, a grammar cannot account for the fact that some sentences which differ ONLY morphemically are interpreted as different in meaning (e.g. *The tiger bit me* and *The mouse bit me*) while other sentences which differ only morphemically are interpreted as identical in meaning (e.g. *The oculist examined me* and *The eye doctor examined me*). Second, a grammar cannot account for the fact that some sentences of radically different syntactic structure are synonymous (e.g. *Two chairs are in the room* and *There are at least two things in the room and each is a chair*) while other syntactically different sentences are not. In each case, the interpretation of the sentences is determined in part by the meanings of their morphemes and by semantic relations among the morphemes. The reason for including a dictionary as a component of a semantic theory is precisely to provide a representation of the semantic characteristics of morphemes necessary to account for the facts about sentences and their interrelations that the grammar leaves unexplained.

What has always been unclear about a semantic theory is what component(s)

it contains besides a dictionary, and how the components of a semantic theory relate to one another and to the grammar. We can find this out by asking in what respects a dictionary and grammar alone are NOT sufficient to match the fluent speaker's interpretations of sentences.

Let us imagine a fluent speaker of English presented with the infinite list of sentences and their structural descriptions generated by a grammar of English. Given an accurate dictionary of English WHICH HE APPLIES BY USING HIS LINGUISTIC ABILITY, the fluent speaker can semantically interpret any sentence on the list under any of its grammatical derivations. He can determine the number and content of the readings of a sentence, tell whether or not a sentence is semantically anomalous, and decide which sentences on the list are paraphrases of each other. Now contrast the fluent speaker's performance with the performance of a machine which MECHANICALLY¹³ applies an English dictionary to a sentence in the list by associating with each morpheme of the sentence its dictionary entry. It is clear that the dictionary usually supplies more senses for a lexical item than it bears in almost any of its occurrences in sentences. But the machine will not be able to select the sense(s) which the morpheme actually bears in a given sentence context, except so far as the selection is already determined by the grammatical markers assigned to the morpheme in the derivation of the sentence. (Thus the machine will be able to choose the correct sense of *seal* in *Seal the letter* so far as the choice is determined by the fact that in this sentence *seal* is marked as a verb, and the correct sense of *seal* in *The seal is on the letter* so far as the choice is determined by the fact that in this sentence *seal* is marked as a noun. But the machine will not be able to distinguish the correct sense of *seal* in *One of the oil seals in my car is leaking* from such incorrect senses as 'a device bearing a design so made that it can impart an impression' or 'an impression made by such a device' or 'the material upon which the impression is made' or 'an ornamental or commemorative stamp' and so forth, since all of these senses can apply to nominal occurrences of *seal*.) What the machine is failing to do is to take account of or utilize the semantic relations between morphemes in a sentence. Hence it cannot determine the correct number and content of readings of a sentence. Nor can it distinguish semantically anomalous sentences from semantically regular ones. Since the machine will associate a dictionary entry with each morpheme in a sentence, it does not distinguish cases in which the sense of a morpheme or string of morphemes in a sentence precludes other morphemes in the sentence from bearing ANY of the senses that the dictionary supplies for them. (E.g. the machine cannot distinguish *The wall is covered with silent paint* from *The wall is covered with fresh paint*.) Finally, the machine cannot tell which sentences in the list are paraphrases of each other in

¹³ The qualification 'mechanically' is important: it precluded the employment of linguistic skills not represented by the grammar or the dictionary. It is precisely the possession of such skills that distinguishes the fluent speaker from the nonspeaker equipped with a grammar and a dictionary. Hence, the degree to which the nonspeaker is permitted access to such skills is the degree to which we obscure what must be accounted for. Conversely, by prohibiting their employment, as we do by the qualification 'mechanically', we bring into clear relief the skills that a semantic theory of a natural language must account for.

any case except the one in which the sentences are of exactly the same syntactic structure and the corresponding words are either identical or synonymous.

The comparison between a fluent speaker and a machine reveals the respects in which a grammar and dictionary by themselves do not suffice to interpret sentences like a speaker of the language. What the fluent speaker has at his disposal that a machine has not are rules for applying the information in the dictionary—rules which take account of semantic relations between morphemes and of the interaction between meaning and syntactic structure in determining the correct semantic interpretation for any of the infinitely many sentences which the grammar generates. Thus, a semantic theory of a natural language must have such rules (which we shall call 'projection rules') as one of its components if it is to match the speaker's interpretations of sentences.

The central problem for such a theory is that a dictionary usually supplies more senses for a lexical item than it bears in an occurrence in a given sentence, for a dictionary entry is a characterization of EVERY sense that a lexical item can bear in ANY sentence. Thus, the effect of the projection rules must be to select the appropriate sense of each lexical item in a sentence in order to provide the correct readings for each distinct grammatical structure of that sentence. The semantic interpretations assigned by the projection rules operating on grammatical and dictionary information must account in the following ways for the speaker's ability to understand sentences: they must mark each semantic ambiguity that a speaker can detect; they must explain the source of the speaker's intuitions of anomaly when a sentence evokes them; they must suitably relate sentences that speakers know to be paraphrases of each other.¹⁴

Pictured in this way a semantic theory interprets the syntactic structure which the grammatical description of a language reveals. This conception gives content to the notion that a semantic theory of a natural language is analogous to a model which interprets a formal system. Further, it explicates the exact sense of the doctrine that the meaning of a sentence is a function of the meanings of its parts. The system of projection rules is just this function.

6. The structure and evaluation of dictionary entries. We shall here describe the form that a dictionary entry must take in a semantic theory, and discuss how, in an empirical study of the semantics of a natural language, we can evaluate the adequacy of proposed dictionary entries for the lexical items of that language. The next section will describe the form of the projection rules.

From the viewpoint of a semantic theory, a dictionary entry consists of two

¹⁴ The distinction between the dictionary and the rules for its application corresponds, in psychological terms, to a difference between mental processes. The dictionary is something that the speaker learns item by item, more or less by rote, and is constantly learning more of. Knowledge of the rules for applying the dictionary, on the other hand, is gained early and in toto, and comes into play whenever a speaker uses his language. Correspondingly, the use of what is learned in learning a dictionary depends on recalling relatively independent bits of information. The rules involve the exercise of a faculty for coding and decoding linguistic information; they organize whatever systematic, nongrammatical information the speaker has about his language and are thus, in the strongest sense, essential to a knowledge of the language. To know a natural language one must know these rules, but one need not know more than a small fraction of its vocabulary.

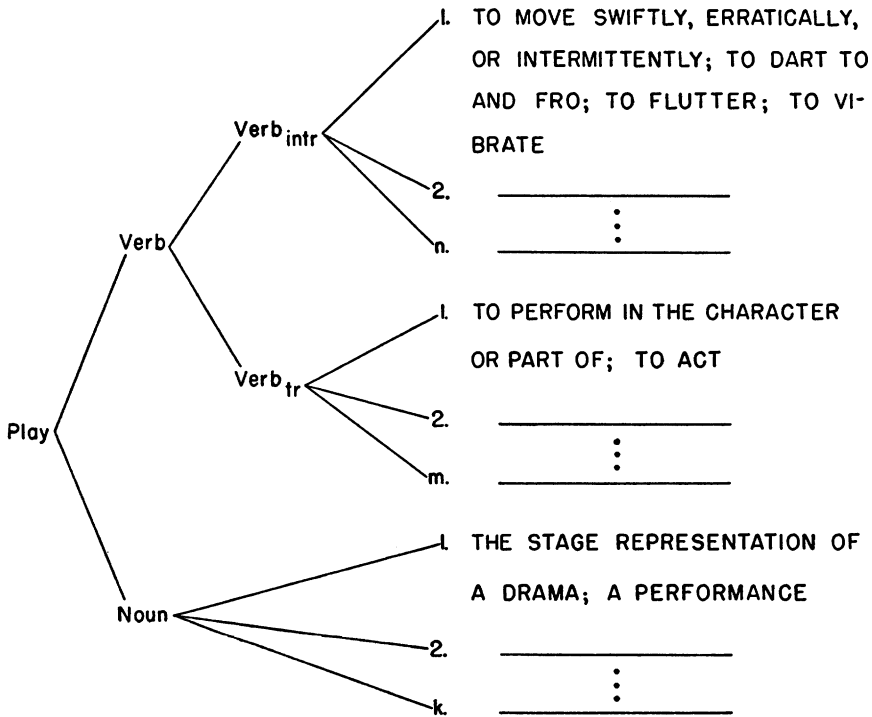


FIG. 2

parts: a grammatical section which provides the part-of-speech classification of the lexical item, and a semantic section which represents each of the distinct senses of the lexical item in its occurrences as a given part of speech. (This leaves out much of what is conventionally found in a dictionary entry, e.g. pronunciation, etymology, chronology. Such information is not relevant to a synchronic semantic description of a language.) For example, the word *play* receives an entry which has grammatical and semantic components as in Fig. 2. The grammatical section classifies the syntactic roles which the lexical item can play in sentences, while the semantic portion supplies one SENSE of the lexical item as the terminal element of each complete distinct descending path through the tree which represents the entry. The sense terminating each path can in turn be analyzed into two parts: a SENSE-CHARACTERIZATION (which appears mandatorily) and a sequence of one or more synonyms (which appears optionally).

The central concept to be studied in this section is that of a sense-characterization of a lexical item. We can justify our concern with this, to the exclusion of synonyms, on the ground that the concept 'synonymity' can be reconstructed in terms of the concept 'sense-characterization' but not conversely. Therefore, the information about synonyms which a dictionary must provide can be given solely in terms of sense-characterizations. In particular, two lexical items have

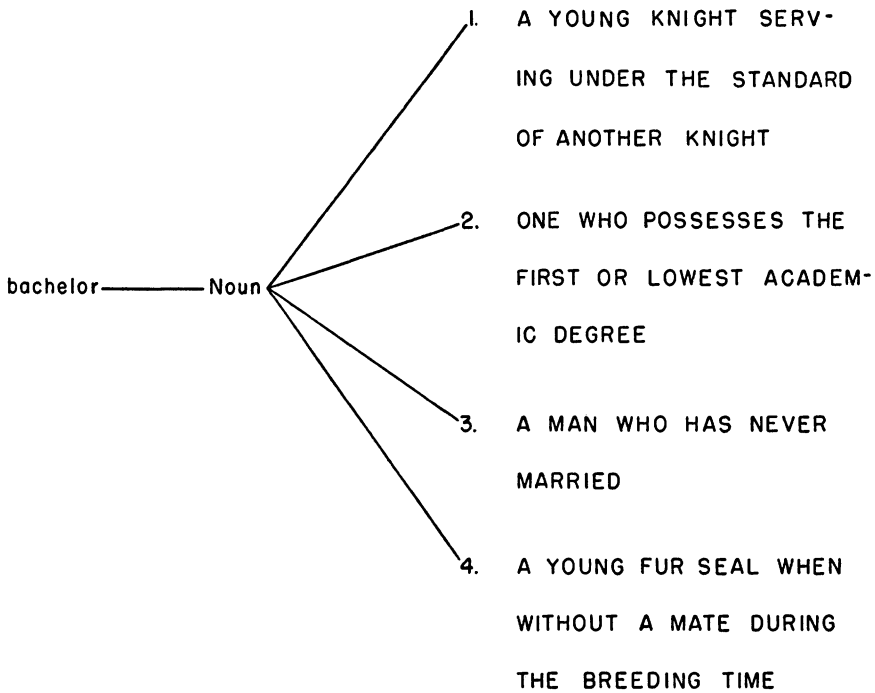


FIG. 3

n synonymous senses if and only if they have n paths in common, and two lexical items are fully synonymous if and only if they have identical entries, i.e. if every path of one is a path of the other. The explicit inclusion of synonyms in a dictionary entry, which is the common practice of conventional dictionaries, is a redundancy introduced to save the user the effort of discovering the synonyms of a lexical item by comparing its sense-characterizations with those of every other item in the dictionary. In short, the practice of listing the synonyms of an item is simply a technique of cross reference. This follows from the fact that it must be a condition upon the adequacy of a dictionary that items which are synonymous in n of their senses have n paths in common.

For the word *bachelor* dictionaries¹⁵ give substantially the entry diagrammed in Fig. 3. However, for reasons which will presently be made clear, the presentation of dictionary entries in the form exemplified in Figs. 2 and 3 is not adequate for a semantic theory. Instead, we require entries in a form exemplified in Fig. 4. Here the unenclosed elements are GRAMMATICAL MARKERS, the elements enclosed in parentheses are what we shall call SEMANTIC MARKERS, and the expressions enclosed in brackets are what we shall call DISTINGUISHERS. We have already commented upon the function of grammatical markers. The semantic markers and distinguishers are the means by which we can decompose the meaning of one

¹⁵ Our sources for dictionary information throughout this paper have been *The shorter Oxford English dictionary* and *Webster's new collegiate dictionary*.

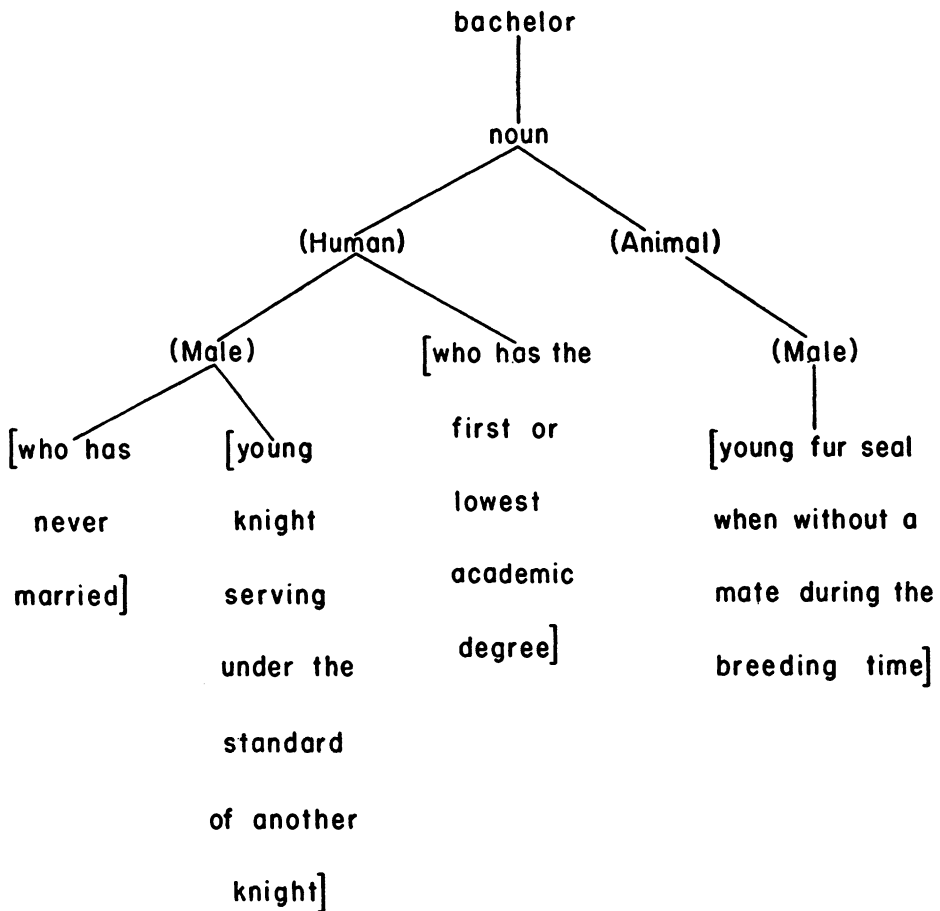


FIG. 4

ense of a lexical item into its atomic concepts, and thus exhibit the semantic structure IN a dictionary entry and the semantic relations BETWEEN dictionary entries. That is, the semantic relations among the various senses of a lexical item and among the various senses of different lexical items are represented by formal relations between markers and distinguishers.

It is clear that any lexical information which a conventional dictionary entry can represent can also be represented by an entry in the normal form shown in Fig. 4. It is also clear that any semantic relations which can be reconstructed from an entry of the former type can also be reconstructed from one of the latter. Distinct senses continue to be represented as distinct paths, synonymous senses of a lexical item continue to be represented in terms of identity of paths, and so on. On the other hand, there are semantic relations which can be reconstructed from entries in our normal form but not from entries in the conventional dictionary form. One such relation is that of SEX-ANTONYMY. This relation holds

between the members of such pairs of words as *bachelor* and *spinster*, *man* and *woman*, *aunt* and *uncle*, *bride* and *groom*, *brother* and *sister*, *cow* and *bull*. What formally characterizes a sex-antonymous pair of words is that the members have identical paths except that where one has the semantic marker (Male) the other has the semantic marker (Female). Since there are indefinitely many important semantic relations which cannot be formally reconstructed from entries in the conventional dictionary, conventional dictionary entries have a serious theoretical disadvantage. But that disadvantage is not the primary reason for introducing our normal form. This is that a formalization of the conventional dictionary entry is required in order to permit a formal statement of the projection rules. We shall go into more detail later.

Semantic markers are the elements in terms of which semantic relations are expressed in a theory. Here there is a strong analogy to grammatical markers, since a grammatical marker (Noun, Verb, Adjective, etc.) is an element in terms of which syntactic relations are expressed. The semantic markers assigned to a lexical item in a dictionary entry are intended to reflect whatever systematic semantic relations hold between that item and the rest of the vocabulary of the language. On the other hand, the distinguishers assigned to a lexical item are intended to reflect what is idiosyncratic about its meaning. Generally speaking, a change in the system of semantic markers has extensive consequences throughout the semantic theory, i.e. such a change radically alters the semantic relations which the theory claims to find between indefinitely many words in the language. But a change in a distinguisher merely alters the relation between one item and its synonyms. For example, if the distinction between the markers (Male) and (Female) were obliterated in a semantic theory of English, not only would every pair of sex-antonyms be represented as synonymous but the indefinitely many other semantic relations involving this distinction would also be incorrectly represented by the theory. In contrast, eliminating the distinguisher [young fur seal when without a mate during the breeding time] would merely prevent a theory from representing one sense of *bachelor* and whatever synonymy relations obtained between that sense of *bachelor* and certain senses of other words.

Branching under a semantic marker is sometimes singular but very often dyadic or greater.¹⁶ Since every path in a dictionary entry represents a distinct sense of a lexical item, a lexical item whose dictionary entry contains polyadic branching has more than one sense, i.e. it is ambiguous. From the viewpoint of the semantic interpretation of sentences, polyadic branching represents the possibility of sentential semantic ambiguity in any sentence in which the ambiguous lexical item appears. For a necessary condition on the semantic ambiguity of a sentence is that it contain an ambiguous lexical item. But clearly this condition is not also sufficient, since not all sentences containing ambiguous lexical items are themselves ambiguous. Consider the sentence *The stuff is light*

¹⁶ In the entries for some lexical items, there will be paths in which the lowest semantic marker dominates nothing, i.e. paths which do not terminate in distinguishers. Such lexical items have special theoretical significance: they are the natural language's representation of semantic categories. Cf. J. J. Katz and J. A. Fodor, 'Categories', unpublished.

enough to carry. The dictionary entry for the word *light* exhibits branching into the semantic markers (Color) and (Weight). Such branching is required to account for the ambiguity of such sentences as *The stuff is light*, *He wears a light suit in the summer*. But since *The stuff is light enough to carry* is unambiguous, it follows that the expression *enough to carry* somehow selects one of the paths in the dictionary entry for *light* and excludes the other(s). Hence, the semantic interpretation of *The stuff is light enough to carry* must explain why the occurrence of *light* in this sentence is understood according to the sense in which *light* is a weight adjective.

In short, if a semantic theory is to predict correctly the number of ways in which speakers will take a sentence to be ambiguous and the precise content of each term of each ambiguity, it must be able to determine every case in which a sentence containing ambiguous lexical items is itself ambiguous and every case in which selection resolves the ambiguities. But this, in turn, amounts to accepting the condition that a dictionary must be so constructed that every case of lexical ambiguity is represented by polyadic branching and that every case of selection can be represented as the exclusion (by some sentence material) of one or more branches. Semantic anomaly can then be construed as the limiting case of selection: the case where there is a lexical item in a sentence whose paths are ALL excluded by selections due to other material in the sentence.

Given the principle that semantic relations are expressed in terms of semantic markers alone, we can see that the primary motivation for representing lexical information by semantic markers will be to permit a theory to express those semantic relations which determine selection and thereby to arrive at the correct set of readings for each sentence. That selection must be represented in terms of semantic markers follows from the fact that selection is a semantic relation between parts of a sentence, together with the principle that all semantic relations are expressed by semantic markers. Thus, the markers in each entry in the dictionary must be sufficient to permit us to reconstruct the operation of the mechanisms of selection in each of the sentences in which the lexical item receiving that entry appears.

Another consequence of expressing semantic relations solely in terms of semantic markers is that distinguishers, when they appear in a path in a dictionary entry, must appear as terminal elements, i.e. there must be no branching under a distinguisher. If branching under a distinguisher were allowed, the theory would posit at least one semantic relation which its dictionary failed to represent by semantic markers, viz. the one between the senses of the lexical item differentiated by that branching.

The distinction between markers and distinguishers is meant to coincide with the distinction between that part of the meaning of a lexical item which is systematic for the language and that part which is not. In order to describe the systematicity in the meaning of a lexical item, it is necessary to have theoretical constructs whose formal interrelations compactly represent this systematicity. The semantic markers are such constructs. The distinguishers, on the other hand, do not enter into theoretical relations within a semantic theory. The part of the

meaning of a lexical item that a dictionary represents by a distinguisher is the part of which a semantic theory offers no general account.

We must now consider the basis on which to decide to represent some lexical information by semantic markers and other lexical information by distinguishers. In the last analysis, the decision can only be justified by showing that it leads to correct interpretation of sentences. What must be explained, therefore, is how such decisions affect the assignment of semantic interpretations and, conversely, how the requirement that a theory assign semantic interpretations correctly affects decisions about the way in which a piece of lexical information is to be represented.

A particular semantic theory of a natural language can REPRESENT only those sentential semantic ambiguities that result from the occurrence of a lexical item for which the dictionary of the theory provides an entry with two or more paths. The degree of semantic ambiguity that a semantic interpretation assigns to a sentence is a function of the degree of branching within the entries for the lexical items appearing in the sentence—branching into markers or into distinguishers, or a combination of both counting equally in determining the degree of ambiguity. On the other hand, a particular semantic theory of a natural language can RESOLVE only those sentential semantic ambiguities which result from the occurrence of lexical material associated with dictionary entries containing two or more paths that differ by at least one semantic marker. This limitation on the power of a semantic theory to resolve ambiguities is a direct consequence of the fact that selection can operate only upon semantic markers. Hence, decisions to represent a piece of lexical information by markers or distinguishers determine in part what semantic ambiguities will be only marked in the semantic interpretation of sentences and which ones will be both marked and resolved.

Such a decision is controlled by two kinds of considerations. Since we wish to construct a semantic theory in such a way that its output matches the performance of a fluent speaker, we want the theory to represent in its semantic interpretations just those ambiguities that the fluent speaker can mark and to resolve just those ambiguities that he can resolve. This will mean that in constructing a theory the lexical information to be represented by markers and by distinguishers will be controlled by our evidence about the disambiguations which a fluent speaker can make. If the dictionary entry for *bachelor* is given as in Fig. 4, every sentence in which this word appears will be represented as ambiguous between the senses given by the paths *bachelor* → noun → (Human) → (Male) → [who has never married], and *bachelor* → noun → (Human) → (Male) → [young knight serving under the standard of another knight]. Since this ambiguity of *bachelor* is represented only by a difference of distinguishers, there is no way that a theory whose dictionary contains this entry can resolve it. But though this is an absolute limitation on such a theory, it is not an absolute limitation on the construction of semantic theories in general. If we notice that fluent speakers do not take such sentences as *The old bachelor finally died* to be ambiguous, we can construct our semantic theory to accommodate this fact simply by taking the lexical information that a bachelor in the second sense is necessarily young to

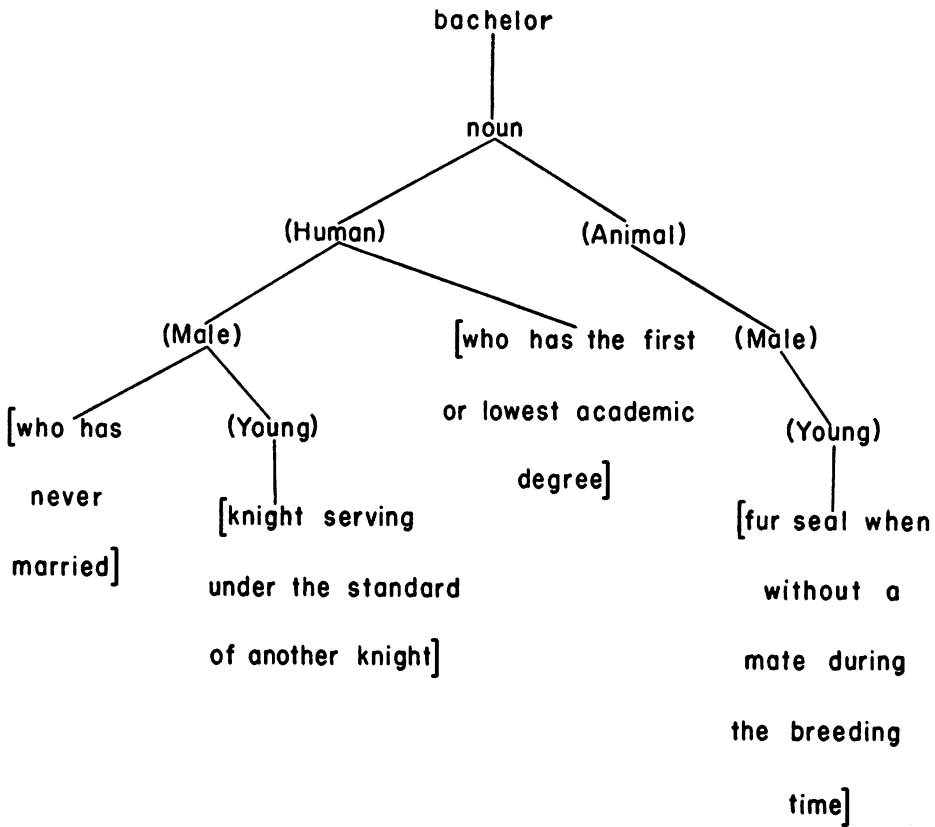


FIG. 5

be marker information rather than distinguisher information. This is done by adding the marker (Young) to the marker system and rewriting the dictionary entry for *bachelor* according to Fig. 5.

The other kind of consideration that controls what lexical information is to be included in the system of semantic markers is the desire for systematic economy. The addition of new semantic markers, as in Fig. 5, is for the sake of increasing the precision and scope of a semantic theory, but in so doing it also increases the complexity of the theory's conceptual apparatus. Since allowing more complexity often coincides with greater precision and scope, the decision should be made on the basis of a strategy which seeks to maximize systematic economy: the greatest possible conceptual economy with the greatest possible explanatory and descriptive power. If such decisions are optimally made, there should eventually come a point when increasing the complexity of a semantic theory by adding new markers no longer yields enough advantage in precision or scope to warrant the increase. At that point, the system of semantic markers should reflect exactly the systematic features of the semantic structure of the language.

So far we have reconstructed four types of information which conventional dictionaries provide about a lexical item: its part-of-speech classification, the number of its senses, its systematic semantic features, and its idiosyncratic features. There is one further type of information in conventional dictionaries that is relevant to synchronic semantic description: the relation between features of certain combinations into which a lexical item enters and the sense which the item bears in those combinations. For example, consider *The shorter Oxford English dictionary's* entry for the word *honest*: '... 3. of persons: of good moral character, virtuous, upright, ... of women: chaste, virtuous, ...' The phrases 'of persons' and 'of women' are intended to indicate that the senses that follow them apply only under the conditions that they specify. That is, these specifications indicate that if the nominal head which *honest* modifies refers to a person without specification of sex, then *honest* has the meaning 'of good moral character, virtuous, upright', and if the nominal head refers to a woman, then *honest* means EITHER 'of good moral character, virtuous, upright' OR 'chaste, virtuous'. Our reconstruction of this type of dictionary information must follow conventional dictionary procedure as far as it goes, but should go further in that the reconstruction should provide ALL the information necessary to determine selection and exclusion. Where the conventional dictionary, by using devices like the phrases with 'of ...', tells us what a word means in certain combinations, our reconstruction must do this systematically and also provide a basis for determining what combinations are semantically acceptable and which ones are not.

For our reconstruction, we shall use left and right angles enclosing a function of syntactic or semantic markers. Such configurations of symbols will be affixed to the terminal element of a path (either the distinguisher or the last semantic marker if there is no distinguisher) and will be construed, relative to the projection rules, as providing a necessary and sufficient condition for a semantically acceptable combination. The angle-enclosed material terminally affixed to the path of a modifier determines the applicability of that path of the modifier to a sense of a nominal head. In particular, a path in the dictionary entry for *honest* will be: *honest* → adjective → (Evaluative) → (Moral) → [innocent of illicit sexual intercourse] <<(Human) & (Female)>>. This is to be construed as saying that an adjectival occurrence of *honest* receives the interpretation (Evaluative) → (Moral) → [innocent of illicit sexual intercourse] just in case the head it modifies has a path containing both the semantic markers (Human) and (Female). How in actual practice a semantic theory utilizes angle-enclosed material to determine selection and exclusion relations in order to obtain semantic interpretations of sentences can only be made clear by the statement of the projection rules.

This concludes the characterization of our normal form for dictionary entries. A dictionary is, then, a list (ordered or not) of the lexical items of the language, each item being associated with an entry in our normal form. The question whether the items are to be words, morphemes, or other units we do not attempt to decide here; but certain considerations are relevant to the decision. The most important is that we choose the unit that will enable us to describe the largest amount of the compositional structure of the language. As a rule, the meaning of a word is a compositional function of the meanings of its parts, and we would

like to be able to capture this compositionality. An approach which directs us to choose as lexical units the most compositionally basic units of the language has, moreover, simplicity in its favor. Wherever we can use composition, dictionary entries are avoided. Thus, instead of having an entry for each verb that takes the prefix *de* and a separate entry for *de* plus that verb, we must choose our lexical units so that the dictionary need only contain an entry for *de* and an entry for the unprefix form of each verb. This economy can be achieved because combinations of *de* + verb are compositional wherever the verb is semantically marked as (Process) → (Reversible).

It will be noticed that the dictionary is so formulated that all semantic properties and relations represented in entries are FORMALLY represented. This is required so that, given a formal statement of the projection rules (i.e. a statement in which the application of rules is defined solely in terms of the shapes of the symbols they apply to, and the operations which the rules effect in producing their output are mechanical), the question what semantic interpretation is assigned to a given sentence can be answered by formal computations without the aid of linguistic intuitions or insights. The need for a formal semantic theory derives from the need to avoid vacuity. A semantic theory is vacuous to the extent that the speaker's intuitions or insights about semantic relations are relied on in order to apply the rules of the theory correctly. Thus, it is uninformative to be told that an English sentence exhibits a semantic relation R just in case it satisfies the condition C, if C is so formulated that we cannot know whether C is satisfied without relying on a speaker's intuitive knowledge of semantic relations like R. A formal theory is ipso facto not vacuous in this respect, since no knowledge about semantic relations in any language is required to determine the correct application of its rules.

Now we turn to the problem of evaluating the adequacy of dictionary entries. It is often assumed that a semantic theory must yield a feasible mechanical procedure which enables the linguist to actually construct a dictionary from information about the verbal behavior of speakers. Every proposal for such a procedure, however, has proved a complete failure; we believe that this is in the nature of the case. We also think that theorists who insist upon a mechanical procedure for deciding whether a putative dictionary entry is optimal have set their aims too high; the practical impossibility of such a decision procedure is also, we believe, in the nature of the case. We shall not argue directly for these claims. We make them primarily to warn the reader against construing the conception of a semantic theory proposed in this paper as either a mechanical discovery procedure or a mechanical decision procedure for dictionary entries.

However, the present paper can be understood as proposing a conception of semantic theory which provides, inter alia, a procedure for determining which of two proposed dictionary entries is the better for a given language. This evaluation procedure, be it noted, differs considerably from what is usually envisioned by semantic theorists. In our conception, a dictionary is only one component of a semantic theory which has as its other component a set of projection rules for semantically interpreting sentences on the basis of the dictionary. Only the theory as a whole can be subjected to empirical test. This means that if a se-

semantic theory gives incorrect interpretations for sentences, one must then decide whether to revise some dictionary entries, some projection rules, or some of each. None the less, there is a derivative sense in which questions of evaluation can be raised about particular dictionary entries. Given projection rules and other dictionary entries that are sufficiently well established, which of two proposed entries yields the best interpretations for sentences? This conception of evaluating dictionary entries differs from the usual one in that it makes evaluation a matter of the degree to which the entry helps achieve the purpose of a dictionary within a theory of semantic interpretation. Semantic theorists usually think of such evaluation as effected by criteria which select the preferable entry simply on the basis of facts about the verbal behavior of speakers, thus overlooking the fact that it is the interpretation of sentences, not the construction of dictionaries, that is the objective of a semantic theory. Because they have overlooked this, their criteria for evaluating dictionary entries are invariably too weak: they fail to utilize systematic constraints on the semantic interpretation of sentences (matching the fluent speaker's ability to determine the number of readings of a sentence, the content of the readings, and their paraphrase relations) in choosing a preferable dictionary entry.

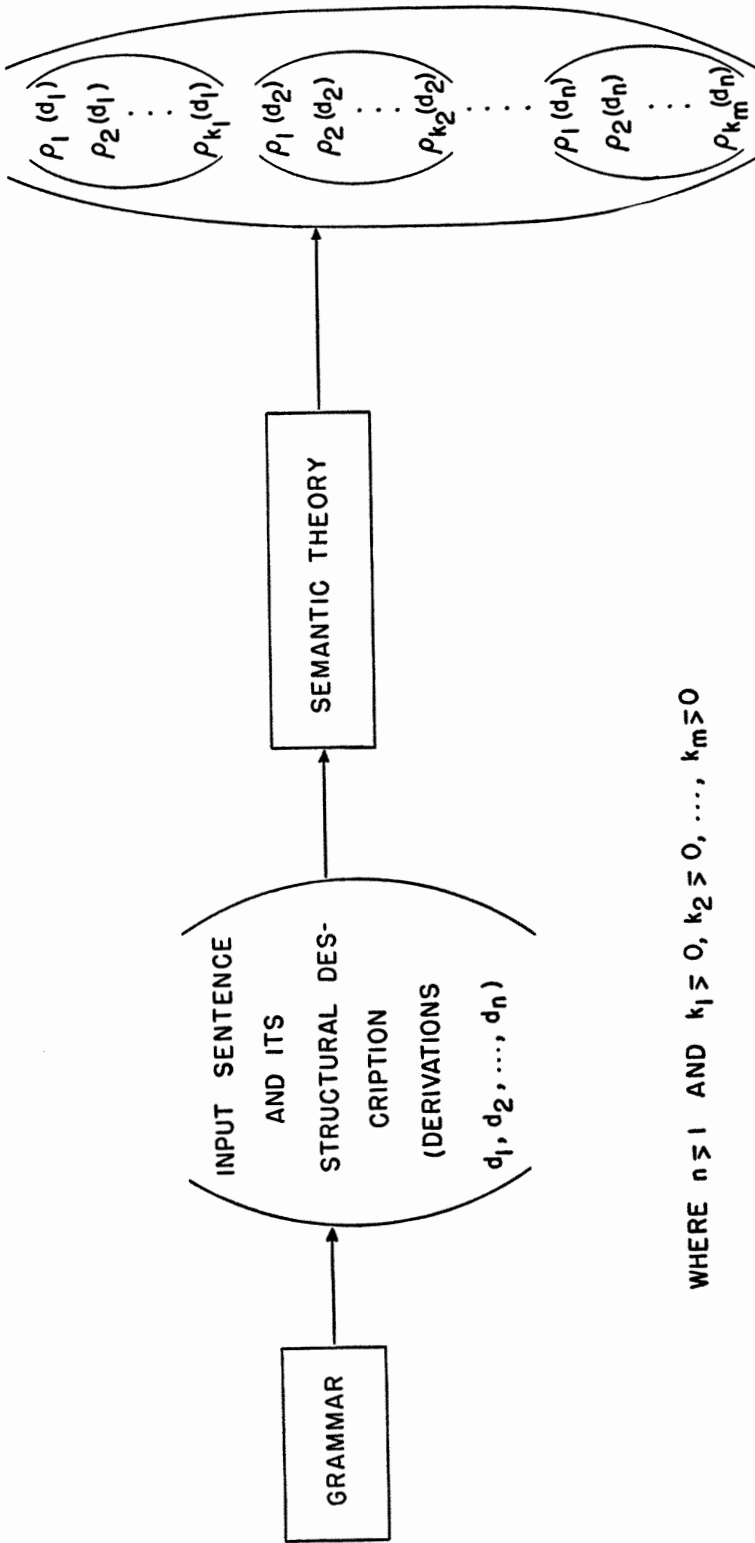
The controls on a semantic theory of a natural language are nothing more than the usual empirical and methodological constraints imposed on any scientific theory: the requirement that a semantic theory match the fluent speaker's ability to interpret sentences is the particular form, in semantics, of the general methodological requirement that a theory accord with the facts. If certain consequences of a semantic theory conflict with the facts (the performance of fluent speakers), various revisions in the dictionary component, in the projection rule component, or in both must be tried out and compared to determine which solution best accommodates the linguistic evidence.

7. The projection rule component. A sentence and its grammatical description provide the input to a semantic theory. Its output is a semantic interpretation of each sentence given as input. We may picture the situation as in Fig. 6.

Fig. 6 shows the input to a semantic theory to be a sentence *S* together with a structural description consisting of the *n* derivations of *S*, *d*₁, *d*₂, ..., *d*_{*n*}, one for each of the *n* ways that *S* is grammatically ambiguous. The output of the semantic theory is shown as *k*₁ readings for *d*₁, *k*₂ readings for *d*₂, ..., *k*_{*m*} readings for *d*_{*n*}, each reading corresponding to a term of a semantic (nongrammatical) ambiguity of *S* on some derivation. The schema $\rho_i(d_j)$ represents the *i*th reading of *d*_{*j*} (which the semantic theory supplies).

We can now characterize the notion 'semantic interpretation of the sentence *S*' as the conjunction $\psi d_1 \ \& \ \psi d_2 \ \& \ \dots \ \& \ \psi d_n$ of the semantic interpretations of the *n* derivations of *S*. The 'semantic interpretation of *S* on the derivation *d*_{*j*}' is the output of the dictionary and projection rule components for *S* on *d*_{*j*} together with the statements about *S* that can be made on the basis of the following conventions:

- (1) If $k_1 + k_2 + \dots + k_m = 1$, then *S* is unambiguous.



WHERE $n \geq 1$ AND $k_1 \geq 0, k_2 \geq 0, \dots, k_m \geq 0$

FIG. 6

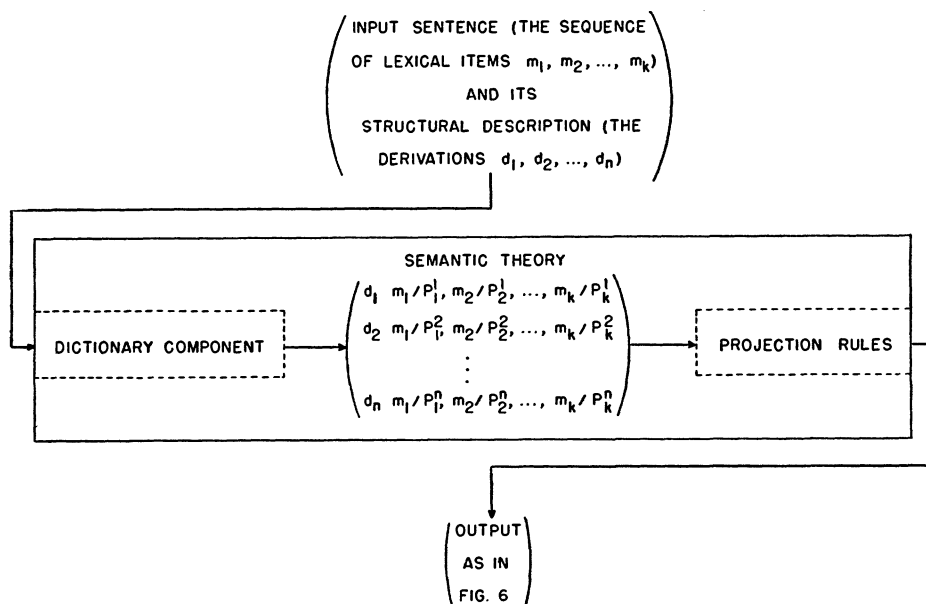


FIG. 7

- (2) If $k_1 + k_2 + \dots + k_m > 1$, then S is $k_1 + k_2 + \dots + k_m$ ways ambiguous.
- (3) If $k_1 + k_2 + \dots + k_m = 0$, then S is fully anomalous (i.e. anomalous on every derivation).
- (4) If the set of readings assigned to the derivation d_j , $\rho_1(d_j), \rho_2(d_j), \dots, \rho_k(d_j)$, has exactly one member, then S is unambiguous on d_j .
- (5) If the set of readings assigned to the derivation d_j has more than one member, then S is k_j ways semantically ambiguous on d_j .
- (6) If the set of readings assigned to d_j is null, then S is semantically anomalous on d_j .
- (7) If S and another sentence P have at least one reading in common, then S and P are paraphrases on that reading.
- (8) If S and P have all readings in common, then S and P are full paraphrases.¹⁷

Fig. 7 schematizes the relation between the dictionary component and the projection rule component. The input to the dictionary component consists of a sentence S represented by a sequence of lexical items m_1, m_2, \dots, m_k and the set of derivations of S . The symbol ' P_j^i ' stands for a finite nonnull set of paths drawn from the dictionary entry for the lexical item m_j in S such that any path in the dictionary entry for m_j is in the set only if the path contains grammatical markers which assign m_j the syntactic role it has on the derivation d_i . The slant line

¹⁷ For further specification of what semantic features of a sentence can be marked in terms of the output of a semantic theory cf. J. J. Katz, 'Analyticity and contradiction in natural language', to be published in 1963.

represents the association between a lexical item and a subset of the set of paths in its dictionary entry. The association is effected by the instruction (I) which, together with the dictionary, comprises the dictionary component:

- (I) For each pair d_i and m_j , the path p in the entry for m_j is assigned to the set P_j^i if and only if p has as its initial subpath the sequence of grammatical markers g_1, g_2, \dots, g_r and the derivation d_i contains the path $g_1 \rightarrow g_2 \rightarrow \dots \rightarrow g_r \rightarrow m_j$.


(I) chooses as relevant to the semantic interpretation of a sentence on a given derivation only those paths from the dictionary entries for each of the lexical items in the sentence which are compatible with the lower-level syntactic structure of the sentence on that derivation. The output of the dictionary component is thus a mapping of a finite nonnull set of paths onto each m_j for each d_i . This output in turn, as Fig. 7 shows, is the input to the projection rules.

We can now give a general picture of the operations whereby the projection rule component converts its input into a semantic interpretation. Each sentence which the grammar makes available for semantic interpretation has associated with it n derivations marking the n ways in which it is structurally ambiguous. Each derivation marks the constituent structure or the derived constituent structure (if the sentence is generated transformationally) in a way that can be represented by a tree diagram. We shall employ such tree diagrams in the following pages, BUT IT IS TO BE UNDERSTOOD THAT PROJECTION RULES CAN ALSO TAKE ACCOUNT OF INFORMATION ABOUT THE TRANSFORMATIONAL HISTORY OF A SENTENCE WHICH IS NOT REPRESENTED IN A TREE DIAGRAM.

Fig. 8 gives the derived constituent structure of the sentence *The man hits the colorful ball*.¹⁸ The dictionary component associates sets of paths with such a tree in the manner specified by (I). Thus, after the application of (I), we have the arrangement shown in Fig. 9. The marking of the lexical items *the*, *man*, *hits*, *the*, *colorful*, and *ball* as Article, Noun concrete, Verb transitive, Article, Adjective, and Noun concrete respectively, which at first glance may seem to have been lost in the application of (I), is actually represented as the common initial subpath of every p in each P_j^i , e.g. P_3^1 is the set of paths all of whose members begin $hits \rightarrow V \rightarrow V_{tr}$.

The general way in which the projection rule component works is by proceed-

¹⁸ It can be argued on grammatical grounds that the phrase *the colorful ball* should be represented simply as:



i.e. without the internal syntactic structure it is given in Fig. 9. This representation does not include the information (required by the projection rules) that *colorful* is a modifier of the head *ball*. But the need for this sort of information does not commit us to the assumption that all branching in derived constituent structure trees is binary. For such information can be obtained by examining the transformational history of the sentence. This is a typical case of the way a projection rule can use information taken from the transformational history of the sentences to which it is applied.

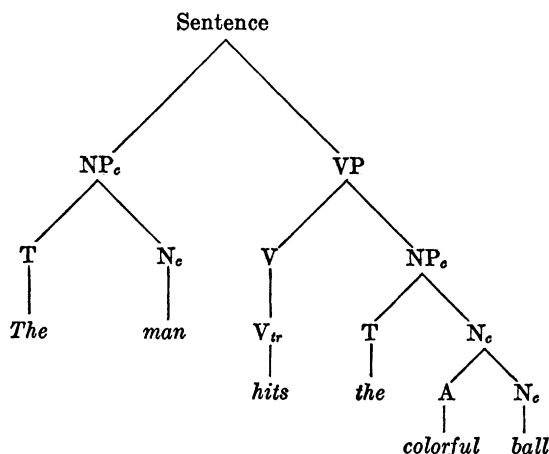


FIG. 8




FIG. 9

ing from the bottom to the top of a constituent structure tree and effecting a series of amalgamations. It starts with the output of (I) and amalgamates sets of paths dominated by a grammatical marker, thus assigning a set of readings to the concatenation of lexical items under that marker by associating the result of the amalgamation with the marker, until it reaches the highest marker 'Sentence' and associates this with a semantic interpretation. The projection rules amalgamate sets of paths dominated by a grammatical marker by combining elements from each of them to form a new set of paths which provides a set of readings for the sequence of lexical items under the grammatical marker. Amalgamation is the joining of elements from different sets of paths under a given grammatical marker if these elements satisfy the appropriate selection restrictions represented by the material in angles.

We now give an example of how a semantic theory of English might interpret a sentence, and in this way exhibit some of the projection rules for English.¹⁹ We

¹⁹ It should be made clear that the rules we shall give are not intended as a contribution to a semantic theory of English but only as examples of the type of rules that such a theory would employ.

choose the sentence *The man hits the colorful ball* under the derivation given in Fig. 8.

The projection rule component receives this sentence and its derivation as input after (I) has operated. (See Fig. 9) The first step for the projection rule component is to amalgamate each set of paths under each of the grammatical markers which immediately dominates ONLY sets of paths, and to associate with the dominating marker the amalgam so obtained. In the case of Fig. 9, the first step is to amalgamate either P_4 and P_5 or P_1 and P_2 ; the order is immaterial.

Let us first take the amalgamation of P_4 and P_5 . The paths comprising the sets P_4 and P_5 are as follows.

P_4

1. *Colorful* → Adjective → (Color) → [Abounding in contrast or variety of bright colors] ((Physical Object) ∨ (Social Activity))
2. *Colorful* → Adjective → (Evaluative) → [Having distinctive character, vividness, or picturesqueness] ((Aesthetic Object) ∨ (Social Activity))

P_5

1. *Ball* → Noun concrete → (Social activity) → (Large) → (Assembly) → [For the purpose of social dancing]
2. *Ball* → Noun concrete → (Physical Object) → [Having globular shape]
3. *Ball* → Noun concrete → (Physical Object) → [Solid missile for projection by an engine of war]

P_{41} is the sense of *colorful* in *The gift came in a colorful wrapper*; P_{42} is the sense in *No novel is less colorful than Middlemarch, excepting Silas Marner*; P_{51} is the sense of *ball* in *The queen danced at the French ambassador's ball*; P_{42} is the sense in *Tennis is played with a ball*; P_{53} is the sense in *The balls whistle free o'er the bright blue sea*. It will be noticed that the sense of *ball* in *He plays ball better than Babe Ruth* is not represented by a path in P_5 , though such a path is to be found in the dictionary entry for *ball*. This is because *ball*, when it means the game, is not a concrete noun; the path which represents that sense is eliminated by (I).

The amalgamation of P_4 and P_5 is accomplished by the following projection rule: (R₁) Given two paths of the form: (1) Lexical String₁ → syntactic markers of head → (a₁) → (a₂) → ... → (a_n) → [1] (set of strings of markers Ω₁); (2) Lexical String₂ → syntactic markers of modifier → (b₁) → (b₂) → ... → (b_m) → [2] (set of strings of markers Ω₂), such that there is a substring σ of the string of syntactic or semantic markers and σ satisfies the condition Ω₂. There is an amalgam of the form: Lexical String₂ + Lexical String₁ → dominating node marker → (a₁) → (a₂) → ... → (a_n) → (b₁) → (b₂) → ... → (b_m) → [[2][1]] (Ω₁), where any (b_i) is null when (∃a_i) (b_i = a_i) and [[2][1]] is [1] when [2] = [1].²⁰

²⁰ The reason why Ω₁ appears in the output of (R₁) is that some heads are, in turn, modifiers of other heads; e.g. adjectives are heads for adverbs and also modifiers for nouns: (*light (red) ball*). In these cases, the conditions in Ω₁ will be required for selection.

The amalgam of P_4 and P_6 is the set of derived paths P_6 :

P_6

1. *Colorful + ball* → Noun concrete → (Social Activity) → (Large) → (Assembly) → (Color) → [[Abounding in contrast or variety of bright colors] [For the purpose of social dancing]]
2. *Colorful + ball* → Noun concrete → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Having globular shape]]
3. *Colorful + ball* → Noun concrete → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Solid missile for projection by an engine of war]]
4. *Colorful + ball* → Noun concrete → (Social Activity) → (Large) → (Assembly) → (Evaluative) → [[Having distinctive character, vividness, or picturesqueness] [For the purpose of social dancing]]

There were six possible amalgamations from the combination of P_4 and P_6 , but only four derived paths because, of the possible combinations, only the combination P_42 and P_61 satisfies the selection restriction $\langle\langle$ (Aesthetic Object) \vee (Social Activity) $\rangle\rangle$. Thus, (R_1) predicts the semantic anomaly of *colorful ball* on the reading where *colorful* has the sense represented by P_42 and *ball* has either the sense represented by P_62 or the sense represented by P_63 .

Another example of how (R_1) contributes to the formalization of the distinction between what is semantically acceptable and what is semantically anomalous is the following. The expression *spinster insecticide* would be regarded as anomalous by speakers of English. This can be predicted on the basis of (R_1) and the dictionary entries for *spinster* and *insecticide*. The relevant path for *spinster* contains: *spinster* → Adjective → (Human) → (Adult) → (Female) → [Who has never married] $\langle\langle$ (Human) $\rangle\rangle$. On the basis of this path, (R_1) assigns no reading to the expression *spinster insecticide*—i.e. predicts that the expression is semantically anomalous—because the path for *insecticide* does not contain the semantic marker (Human) which is necessary to satisfy the selection restriction associated with *spinster*.

(R_1) introduces the semantic markers in the path of the modifier just below the string of semantic markers in the path of the head, eliminating from the path of the modifier all semantic material already present in the path of the head, and associating the distinguishers with one another. The operation of (R_1) corresponds closely to our intuitive notions of the nature of attribution. Attribution is the process of creating a new semantic unit compounded from a modifier and a head, whose semantic properties are those of the head, except that the meaning of the compound is made more determinate than that of the head alone by the information which the compound obtains from the modifier. As Lees comments:²¹

We cannot get along with a single common noun to refer to a familiar common object, but must have at every moment modifiers with which to construct new more complex names

²¹ R. B. Lees, *The grammar of English nominalizations xvii-xviii (IJAL 26:3, 1960)*.

to use for all the specific instances of that object which we encounter and talk about. Thus, we cannot, without extensive ambiguity, refer on every occasion to our favorite beverage by means of the single word *coffee*; instead we name its individual instances with such phrases as 'my coffee', 'that cold cup of coffee you left there', 'some fresh coffee on the shelf', 'a new brand of coffee', 'pretty tasteless coffee', 'Turkish coffee', etc. There is no known limitation on the number of distinct objects for which we must at some time or other have distinctive names, and clearly no dictionary is large enough to contain them all, for a great many of the names which we employ have never before been uttered. Like full sentences themselves, there is no longest name, and there must consequently be an infinity of new names available for us to use when and if the need arises.

Though Lees is commenting on the grammar of nominal compounds, what he says applies equally well to their semantics and to the semantics of other modifier-head constructions. It is only because there is a systematic way of understanding the meaning of such constructions in terms of the meanings of their parts, that the infinite stock of strings produced by the grammatical mechanism for creating new modifier-head constructions can be employed by speakers to refer to familiar objects.

As we have just mentioned, the meaning of a compound is more determinate than the meaning of its head alone in respect of the information which the compound obtains from its modifier(s). The word *aunt* is indeterminate as to age (i.e. both the sentences *My aunt is an infant* and *My aunt is aged* are semantically acceptable), but *spinster*, as we have observed above, contains the semantic marker (Adult) in its path. This marker is carried over to the compound when (R_1) operates to produce an interpretation for *spinster aunt*. Thus, *spinster aunt* is made more determinate (with respect to age) than is *aunt*. This shows up in a comparison between the sentences *My spinster aunt is an infant* and *My spinster aunt is aged*, of which the former is contradictory while the latter is not.

The limiting case, where the addition to the compound of semantic material from the modifier is zero, is of considerable theoretical significance. The compound *unmarried bachelor* is a case in point. The erasure clause in (R_1), 'any b_i is null when $(\exists a_i)(b_i = a_i)$ and $[[2][1]$ is $[1]$ when $[2] = [1]$ ', tells us to delete from the path of the modifier any semantic material already represented in the path of the head. Thus, in forming the compound *unmarried bachelor* all the semantic information in the path of the modifier *unmarried* will be deleted so that the derived path for *unmarried bachelor* will contain no more than the semantic material which comes from the path for *bachelor*. The failure of the modifier to add semantic information would appear to account for the intuition that such expressions as *unmarried bachelor* are redundant and that, correspondingly, such statements as *Bachelors are unmarried* are 'empty', 'tautological', 'vacuous', 'uninformative'. This provides a new explanation of the analyticity of a classical type of analytic truth.²² Moreover, this feature of the projection rules provides another empirical constraint on a semantic theory: if the theory characterizes an expression or sentence as redundant in the above sense, the theory is confirmed if speakers take the expression or sentence in the appropriate way, and is disconfirmed if they do not.

²² Cf. J. J. Katz, 'Analyticity and contradiction in natural language', to appear in *Readings in the philosophy of language*.




FIG. 10

The next step in the semantic interpretation of *The man hits the colorful ball* is the amalgamation of P_1 and P_2 . The entry for *the* in standard dictionaries is exceedingly complex, primarily because the information required to make the correct selections among the various senses of *the* for its sentential occurrences is extremely complicated. We shall have to simplify.

P_1 contains only the path *the* → Noun phrase concrete → Definite Article → [Some contextually definite]. Other paths in the dictionary entry for *the* (those corresponding to the generic senses of the definite article) are not assigned to P_1 by (I) because only the above path contains as its initial subpath the sequence of grammatical markers which dominates *the* in the derivation in Fig. 8.²³ P_2 contains only the path *man* → noun concrete → noun masculine → (Physical Object) → (Human) → (Adult) → (Male). Other paths in the dictionary entry for *man* (those corresponding to the sense of *man* in *Man is occasionally rational* and to the sense in *Every man on board ship was saved except an elderly couple*) do not appear in P_2 , the former because in that sense *man* is not a concrete noun and the latter because in that sense *man* is not a masculine noun. The rule which amalgamates P_1 and P_2 is: (R₂) Given two paths of the form: (1) Lexical String₁ → syntactic markers of noun → semantic markers of head → [1], (2) Lexical String₂ → syntactic markers of article → semantic markers of article → [2] (set of strings of markers Ω), such that there is a substring σ of the string of syntactic or semantic nominal markers and σ satisfies the condition Ω . There is an amalgam of the form: Lexical String₂ + Lexical String₁ → dominating node marker → semantic markers of article → [2] → semantic markers of noun → [1]. The application of (R₂) to P_1 and P_2 produces the derived path: *the* + *man* → Noun phrase concrete → [Some contextually definite] → (Physical Object) → (Human) → (Adult) → (Male). This path is the only member of the set P_7 shown in Fig. 10.

The amalgamation of P_1 and P_6 works in exactly the same way to yield P_3 :

P_3

1. *The* + *colorful* + *ball* → Noun phrase concrete → [Some contextually definite] → (Social Activity) → (Large) → (Assembly) →

²³ In taking NP_c as part of the sequence of grammatical markers in the dictionary entry for *the*, we are not claiming that *the* is a concrete noun phrase, but only that it occurs as an element of a concrete noun phrase and that, when it does, it has the sense in P_1 . This constitutes an extension of the notion of a 'part of speech' classification, but a natural and necessary one.




FIG. 11

- (Color) → [[Abounding in contrast or variety of bright colors] [For the purpose of social dancing]]
2. *The + colorful + ball* → Noun phrase concrete → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast of variety of bright colors] [Having globular shape]]
 3. *The + colorful + ball* → Noun phrase concrete → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Solid missile for projection by an engine of war]]
 4. *The + colorful + ball* → Noun phrase concrete → [Some contextually definite] → (Social Activity) → (Large) → (Assembly) → (Evaluative) → [[Having distinctive character, vividness, or picturesqueness] [For the purpose of social dancing]]

This leaves us with only that part of the constituent structure tree shown in Fig. 11 still to be interpreted.

P_3 is as follows.

1. *hits* → Verb → Verb transitive → (Action) → (Instancy) → (Intensity) → [Collides with an impact] ⟨Subject: (Higher Animal) v (Improper Part) v (Physical Object), Object: (Physical Object)⟩²⁴
2. *hits* → Verb → Verb transitive → (Action) → (Instancy) → (Intensity) → [Strikes with a blow or missile] ⟨Subject: (Human) v (Higher Animal), Object: (Physical Object), Instrumental: (Physical Object)⟩

P_{31} is the sense of *hits* in *The rock hits the ground with a thud*, P_{32} is the sense in *The man hits the ground with a rock*. It will be noticed that the representation of verbs includes between angles selection restrictions upon the subject, objects, and instrumental of the verb. This information is represented by markers of the form 'Subject: α ', 'Object: β ', and 'Instrumental: γ ', where α , β , and γ represent conditions on the paths associated with subject, objects, and instrumentals respectively.

²⁴ Here some explanation is called for. (Instancy) is assigned to verbs representing durationless events. Any sentence whose main verb is marked (Instancy) which is of the form *Subject + Verb + ed + Object + for + numerical quantifier + measure of time* will be understood to mean that the object was verbed more than once. Compare *He hit the ball for three hours* with *He studied the book for three hours*. Next, (Intensity) is assigned to verbs taking adverbs like *hard, soft, gently*. Finally, the marker (Improper Part) is assigned to lexical items that represent wholes which the language contrasts with their parts. The term 'Improper Part' is borrowed from James Thomson.




FIG. 12

A few comments on this dictionary entry for *hits* as a transitive verb. We claim no more than rough accuracy for the characterization given. Our interest here, as throughout the paper, is in prescribing the abstract form of a semantic theory, rather than in actually writing one. The characterization of *hits* is intended primarily to illustrate how the results of a linguistic analysis are to be formally presented so that the projection rules can utilize them. But we have tried to make our examples account for the fundamental semantic features. The failure to mark an achievement sense of *hits* is not an oversight. We choose not to mark the special sense of *hits* as an achievement verb because the behavior of *hits* diverges significantly from that of such paradigmatic achievement verbs as *sees* and *hears*. Thus, unlike *He hit the ball intentionally*, *He saw the picture intentionally* is anomalous (except where it means that he went to see the picture intentionally) and *He heard the music intentionally* is anomalous (except where it means that he didn't just overhear the music). This is perhaps related to the fact that one can intentionally miss the ball, though one cannot in the relevant sense intentionally fail to hear the music. If, however, it should turn out that *hits* must be given a special achievement sense, such a sense can be represented within the formalism of the present paper in a straightforward manner.

The projection rule which amalgamates P_7 and P_8 is: (R_3) Given two paths of the form: (1) Lexical String₁ → syntactic markers of main verb → semantic markers → [1] ⟨sets of strings of markers α , β ⟩, (2) Lexical String₂ → syntactic markers of object of main verb → Remainder of object path, such that there is a substring of the string of syntactic or semantic markers of the object σ and σ satisfies the condition β , there is an amalgam of the form: Lexical String₁ + Lexical String₂ → dominating node marker → semantic markers of main verb → [1] → String analyzed 'Remainder of object path' ⟨set of strings of markers α ⟩. The application of (R_3) to P_7 and P_8 yields P_9 , shown in Fig. 12. P_9 contains the following paths:

P_9

1. *hits* + *the* + *colorful* + *ball* → VP → (Action) → (Instancy) → (Intensity) → [Collides with an impact] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Having globular shape]] ⟨Subject: (Higher Animal) v (Improper Part) v (Physical Object)⟩
2. *hits* + *the* + *colorful* + *ball* → VP → (Action) → (Instancy) → (Intensity) → [Collides with an impact] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Solid missile for projection by engine of war]] ⟨Subject: (Higher Animal) v (Improper Part) v (Physical Object)⟩

3. *hits + the + colorful + ball* → VP → (Action) → (Instancy) → (Intensity) → [Strikes with a blow or missile] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Having globular shape]] ⟨Subject: (Human) v (Higher Animal)⟩
4. *hits + the + colorful + ball* → VP → (Action) → (Instancy) → (Intensity) → [Strikes with a blow or missile] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Solid missile for projection by engine of war]] ⟨Subject: (Human) v (Higher Animal)⟩

Finally, the projection rule which operates on P_7 and P_9 to assign a set of readings to 'Sentence' is: (R_4) Given two paths of the form: (1) Lexical String₁ → syntactic markers of verb phrase → Remainder of verb phrase path, (2) Lexical String₂ → syntactic markers of subject → Remainder of subject path, such that there is a substring σ of the string of syntactic or semantic markers of the subject and σ satisfies the condition α . There is an amalgam of the form: Lexical String₂ + Lexical String₁ → dominating node marker → String analyzed 'Remainder of subject path' → String analyzed 'Remainder of verb phrase path' deleting substring $\langle\alpha\rangle$. The application of (R_4) to P_7 and P_9 yields the set P_{10} :

P_{10}

1. *The + man + hits + the + colorful + ball* → Sentence → [Some contextually definite] → (Physical Object) → (Human) → (Adult) → (Male) → (Action) → (Instancy) → (Intensity) → [Collides with an impact] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Having globular shape]]
2. *The + man + hits + the + colorful + ball* → Sentence → [Some contextually definite] → (Physical Object) → (Human) → (Adult) → (Male) → (Action) → (Instancy) → (Intensity) → [Collides with an impact] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Solid missile for projection by an engine of war]]
3. *The + man + hits + the + colorful + ball* → Sentence → [Some contextually definite] → (Physical Object) → (Human) → (Adult) → (Male) → (Action) → (Instancy) → (Intensity) → [Strikes with a blow or missile] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Having globular shape]]
4. *The + man + hits + the + colorful + ball* → Sentence → [Some contextually definite] → (Physical Object) → (Human) → (Adult) → (Male) → (Action) → (Instancy) → (Intensity) → [Strikes with a blow or missile] → [Some contextually definite] → (Physical Object) → (Color) → [[Abounding in contrast or variety of bright colors] [Solid missile for projection by an engine of war]]

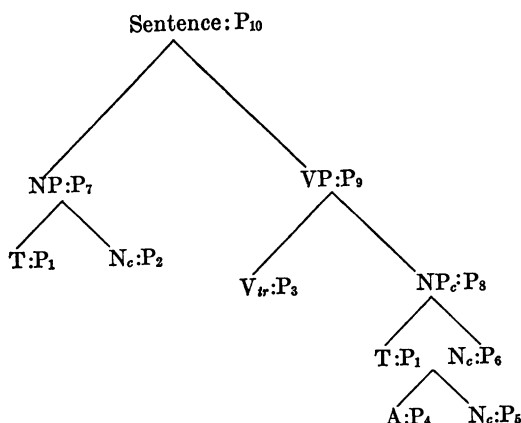


FIG. 13

Accordingly, a semantic theory of English containing rules and entries as given above characterizes the sentence *The man hits the colorful ball* as having the following semantic interpretation. The sentence is not semantically anomalous; it is four ways semantically ambiguous on the derivation in Fig. 8; each term corresponds to a reading in P_{10} ; it is a paraphrase of any sentence which has one of the readings in P_{10} ; and it is a full paraphrase of any sentence that has the set of readings P_{10} assigned to it. The semantic theory interprets the constituent structure tree in Fig. 8 in the way shown in Fig. 13, thus displaying which of the possible combinations of paths at a given node yield derived paths for that node and which possible combinations are blocked.

This completes our example of how a semantic theory of English might interpret a sentence generated by the grammar. Before we conclude our discussion of projection rules, we must consider the question whether the projection rule component will contain types of projection rules different from the type employed above.

What is characteristic of the rules (R_1) through (R_4) is that each rule operates on a part of a partially semantically interpreted constituent structure characterization, amalgamates paths from two sets of paths that are dominated by a particular node, and assigns to that node the set of amalgams as readings for the lexical string that the node dominates. Let us call such rules 'type 1 projection rules'. These rules must assign semantic interpretations to SOME of the sentences generated by the grammar, but they need not be the means by which EVERY sentence receives a semantic interpretation. We can conceive of another type of projection rule (call them 'type 2 projection rules') in the following way. We restrict the application of type 1 projection rules to some formally determined proper subset of the set of sentences. Then we introduce type 2 projection rules to provide a semantic interpretation for every sentence that does not receive a semantic interpretation on the basis of type 1 projection rules. Since those sen-

tences that the grammar produces without the aid of optional transformations, i.e. the kernel sentences, will be semantically interpreted by type 1 projection rules, the type 2 projection rules will assign semantic interpretations to sentences that are constructed with the use of optional transformations. Suppose S has been constructed from a certain set of source sentences by the optional transformation T. A type 2 rule is a rule which operates on the semantic interpretations of these source sentences and on either the derived constituent structure characterization of S or on the transformation T in order to produce a semantic interpretation of S. Type 2 projection rules should assign semantic interpretations in such a way as to reconstruct the manner in which the meaning of a sentence constructed by the transformation T is a function of the meanings of each of the sentences used by T in S's construction.

The basic theoretical question that remains open here is just what proper subsets of the set of sentences are semantically interpreted using type 1 projection rules only. One striking fact about transformations is that a great many of them (perhaps all) produce sentences that are identical in meaning with the sentence(s) out of which the transform was built.²⁵ In such cases, the semantic interpretation of the transformationally constructed sentence must be identical to the semantic interpretation(s) of the source sentence(s), at least with respect to the readings assigned at the sentence level. For example, sentences that are related to each other by the passive transformation, e.g. *The man eats cake* and *Cake is eaten by the man*, have the same meaning, except perhaps in instances where quantifiers are involved.²⁶ Likewise sentence conjunctions, e.g. *The man ate the cake and candy*, which comes from *The man ate the cake* and *The man ate the candy*. Or again, stylistic variants such as *There is something about it that puzzles me*, *There is about it something that puzzles me*, and *There is something that puzzles me about it*. It would be theoretically most satisfying if we could take the position that transformations never change meaning. But this generalization is contradicted by the question transformations, the imperative transformation, the negation transformation, and others. Such troublesome cases may be troublesome only because we now formulate these transformations inadequately, or they may represent a real departure from the generalization that meaning is invariant under grammatical transformations. Until we can determine whether any transformations change meaning, and if some do, which do and which do not, we shall not know what sentences should be semantically interpreted with type 2 projection rules and how to formulate such rules.

Nevertheless, we can decide the cases that are clear. The set of sentences that will be semantically interpreted using type 1 projection rules includes the sentences produced without the aid of optional transformations. Suppose we permit no type 2 projection rule for any transformation that we know preserves meaning, and instead introduce the convention that any sentences related by such a

²⁵ For the background of this point cf. Fodor, 'Projection and paraphrase in semantics', *Analysis* 21:4, 73-7 (1961), and Katz, 'A reply to "Projection and Paraphrase in Semantics"', *Analysis* 22:2.36-41 (1961).

²⁶ In these instances too, if both active and passive have the same meaning because both are ambiguous.

transformation T belong to an equivalence class all of whose members receive the same semantic interpretation. Then the facts that there will always be a kernel sentence in such an equivalence class and that every kernel sentence has a semantic interpretation on the basis of type 1 projection rules mean that every non-kernel sentence in such an equivalence class automatically receives the semantic interpretation of its kernel co-member, which makes them all paraphrases as desired.

This treatment is by far the best method of marking paraphrase relations (and other semantic properties) among stylistic variants which result from the operation of a permutation transformation. This method avoids having a special type 2 rule in each such case: such special type 2 rules have no function except to state the empty fact that these transformations do not affect meaning. This method also avoids the use of type 1 rules on a sentence that is produced by a permutation transformation. This is very desirable, because such transformations produce sentences with derived constituent structure characterizations having far less labelled bracketing than the constituent structure characterization of the source sentence, so that what labelled bracketing survives is generally too little for type 1 rules to be able to interpret semantically the derived sentence.

This treatment has the same merits for transformations that permute so as to produce discontinuous elements in the transform and for transformations that delete material. Thus with this treatment we can most simply account for the paraphrase relations between (and other semantic properties of) such pairs as: *John looked up the rule* and *John looked the rule up*; *Harry plays chess as well as Bill plays chess* and *Harry plays chess as well as Bill*.

The possibility of type 2 projection rules presents two options for the construction of the projection rule component of a semantic theory. Either the projection rule component will consist of type 1 rules alone, or it will contain rules of both type 1 and type 2. Whether type 2 rules will be required, and (if so) to what extent, is a question to which no answer is at present known. The answer involves many considerations, both methodological (conceptual economy, descriptive and explanatory power, etc.) and particular (concerning the structure of individual languages, for instance the degree to which semantic relations between sentences correspond to transformational relations).²⁷

8. Metatheory. We shall here discuss the theoretical perspective from which we have been treating the problem of characterizing the abstract form of a semantic theory—the nature of semantic metatheory. We shall also consider some of the consequences of adopting an explicit metatheory in semantics.

²⁷ We can highlight the difference between theories containing only type 1 rules and theories containing rules of both type 1 and type 2 by contrasting the ways that they would deal with relative clause constructions. A theory containing only type 1 rules would determine from the transformational history of a sentence containing such a construction which nominal the relative suppresses, and would then treat the relative clause as an adjectival on that nominal, with amalgamation proceeding in the normal way. A theory containing both types of rules would first provide an interpretation for the matrix sentence and the embedded sentence which underlie the sentence containing the relative clause, and would then convert the semantic interpretations of the source sentences into a semantic interpretation of the sentence containing the relative.

There are two motives for constructing an explicit metatheory for an area in linguistics, and thus for semantics.²⁸ First, the same scientific curiosity which makes us inquire into the semantic structure of individual languages a fortiori makes us interested in what is common to the semantic structure of families of language or of all languages. Hence a metatheory for semantics must be a theory which represents semantic universals. Second, there must be WELL-ESTABLISHED criteria for choosing among different semantic theories for the same language, where each theory is, as far as we can tell, compatible with the available evidence from fluent native speakers. But if a set of such criteria is to be well-established, it must itself be shown to give desirable results with a wide variety of different languages, i.e. it must choose the better semantic theory over the worse consistently from language to language. Hence a semantic metatheory must provide criteria for evaluating individual semantic theories and establish the adequacy of such criteria. We can satisfy both motives if we construct a metatheory which contains an enumeration of the semantic markers from which the theoretical vocabulary of each particular semantic theory is drawn and a specification of the form of the dictionary entries and rules for a semantic theory of a natural language. For the enumeration and the specification provide both a representation of semantic universals and a basis on which to evaluate particular semantic theories. For example, we may adopt the rule that the preferable theory is the one which is rated highest by a metric which compares dictionary entries in the specified form and chooses the theory requiring the smallest number of markers from the enumeration given in the metatheory.

The semantic markers which we have used in our discussions of dictionary entries and projection rules are, of course, only examples. But if we imagine them functioning in a putative semantic theory of English, then the claim for them would have to be that they are drawn from the enumeration of markers provided by the metatheory, just as the claim for the projection rules would have to be that they are in a form specified by the metatheory. In other words, a semantic marker is a theoretical construct which receives its interpretation the semantic metatheory; it is on a par with such scientific constructs as the atom, the gene, valence, and the noun phrase. A marker like (Human) or (Color) is, then, not an English word, but a construct represented by one.

A metatheory for semantics must also exhibit the relations between semantics and other areas of linguistics. We have discussed the relation between grammatical and semantic rules at some length. We now consider the relation between grammatical and semantic markers.

Much confusion has been generated in the study of language by the search for a line between grammar and semantics. This is because students of language who have tried to draw such a line have sought a criterion to determine when a concept expressing something about the structure of a language is syntactic and when it is semantic. But the trouble is that every such criterion seems to

²⁸ The conception of a metatheory for semantics which is sketched below is adapted from Chomsky's conception of a metatheory for grammar, which he refers to as 'linguistic theory'. Cf. *Syntactic structures*, and *The logical structure of linguistic theory* (microfilm).

be invalidated by concepts which can be regarded, apparently with equal justice, as either syntactic or semantic. There appears to be an overlap between the sets of syntactic and semantic markers; the markers male, female, human, animal, animate, concrete, and abstract appear to fall into this overlap. But the confusion engendered in the search for a line between grammar and semantics is unwarranted: the overlap exists in name only.

This becomes clear once we stop searching for a criterion to decide which markers are properly syntactic and which semantic, and ask instead whether the line between grammatical and semantic markers can be drawn in terms of the theoretical functions they perform.²⁹ For example, in the grammar the distinction between abstract and concrete nouns is drawn in order to construct adequate rules for generating sentences containing nominalizations. According to Lees,³⁰

... there are certain restrictions on subject/predicate-nominal combinations based on abstractness (as well as perhaps on other lower-order nominal categories). There is a small class of (abstract) nouns which may appear in copula sentences opposite both nominalizations and concrete nominals: The problem is that he went there., The problem is his going there., The problem is his tonsils., etc. for such nouns N_a as *problem, trouble, thing, reason, cause, question*, etc. Nominalizations occur opposite only these latter nouns, while concrete nominals N_c occur opposite either other concretes or one of these latter abstract noun N_a ; That he came home is the trouble., but not *That he came home is that she left., or again: His stomach is the cause., His stomach is an organ., but not: *His stomach is his having gone there.

The distinction between mass and count nouns, analogously, is drawn in order to handle the syntactic relations between nouns and their articles and quantifiers. Thus, the mass noun *blood* in the singular takes *the* and *some* but not numerical quantifiers: *The blood was found* but not **One blood was found*; the distinction between animate and inanimate nouns, and between masculine and feminine nouns, has to do (among other things) with pronoun agreement, e.g. *The girl gave her own dress away*, but not **The girl gave his own dress away* or **The girl gave its own dress away*.

On the other hand, semantic markers are introduced to specify something about the meaning of lexical items. Where it appears that a marker is common to both grammar and semantics, there are two distinct markers with the same or similar names. This is most clear from the fact that it is often NOT the case that a lexical item receiving a certain grammatical marker also receives the corresponding semantic marker. If we always assigned a semantic marker where the corresponding grammatical marker is assigned, many lexical items will be given the wrong sense characterization. Grammatically the words *ship, England, fortune*, and *fate* are marked feminine, but clearly they cannot receive the se-

²⁹ It is not at all clear that the request for such a criterion is a reasonable one. Would one ask for an analogous criterion to distinguish the concepts of physics from those of chemistry?

³⁰ Lees, *The grammar of English nominalizations* 14.

semantic marker (Female) if sentences are to receive the correct semantic interpretations. Again, the words *pain*, *ache*, *twinge*, etc. must be marked as concrete nouns, but they cannot be marked as (Physical Object) if we are to account for such anomalies as *The pain weighs three pounds*. Conversely, if we always assigned a grammatical marker where the corresponding semantic marker is assigned, either the grammar will fail to generate some grammatical sentences or it will generate some ungrammatical strings, or else it will fail to assign structure properly. Semantically the nouns *child*, *baby*, and *infant* must be marked as (Human) to obtain correct sense characterizations and correct semantic interpretations; but if they are marked as human nouns, the grammar will fail to generate such sentences as *The baby lost its rattle*.

Grammatical and semantic markers have, then, different theoretical import. Grammatical markers mark the formal differences on which the distinction between well-formed and ill-formed strings of morphemes rests, while semantic markers give each well-formed string the conceptual content that permits it to be a means of genuine verbal communication. They are concerned with different kinds of selection and they express different aspects of the structure of a language. We can justifiably regard semantic markers as theoretical constructs distinct from the markers employed in grammatical description.

LINKED CITATIONS

- Page 1 of 1 -



You have printed the following article:

The Structure of a Semantic Theory

Jerrold J. Katz; Jerry A. Fodor

Language, Vol. 39, No. 2. (Apr. - Jun., 1963), pp. 170-210.

Stable URL:

<http://links.jstor.org/sici?sici=0097-8507%28196304%2F06%2939%3A2%3C170%3ATSOAST%3E2.0.CO%3B2-J>

This article references the following linked citations. If you are trying to access articles from an off-campus location, you may be required to first logon via your library web site to access JSTOR. Please visit your library's website or contact a librarian to learn about options for remote access to JSTOR.

[Footnotes]

¹² **Discourse Analysis**

Zellig S. Harris

Language, Vol. 28, No. 1. (Jan. - Mar., 1952), pp. 1-30.

Stable URL:

<http://links.jstor.org/sici?sici=0097-8507%28195201%2F03%2928%3A1%3C1%3ADA%3E2.0.CO%3B2-4>

²⁶ **Projection and Paraphrase in Semantics**

J. A. Fodor

Analysis, Vol. 21, No. 4. (Mar., 1961), pp. 73-77.

Stable URL:

<http://links.jstor.org/sici?sici=0003-2638%28196103%2921%3A4%3C73%3APAPIS%3E2.0.CO%3B2-O>

²⁶ **A Reply to "Projection and Paraphrase in Semantics"**

Jerrold J. Katz

Analysis, Vol. 22, No. 2. (Dec., 1961), pp. 36-41.

Stable URL:

<http://links.jstor.org/sici?sici=0003-2638%28196112%2922%3A2%3C36%3AART%22AP%3E2.0.CO%3B2-6>

NOTE: *The reference numbering from the original has been maintained in this citation list.*