

Jerry A. Fodor

## HOW TO LEARN TO TALK: SOME SIMPLE WAYS

I should like to say something about what kind of problem the problem of syntax-learning is and something about the kind of theory that might offer a basis for the solution of the problem. I wish to emphasize at the outset, however, that these remarks are very tentative and that, even if they should prove correct, they could not possibly prove exhaustive. To characterize the form of a theory of syntax-learning is to do considerably less than to provide such a theory, and the more we learn about language the harder language-learning looks.

There are, nevertheless, some points that are presumably not in dispute. One is that speaking a language requires information about the structural relations within and among the sentences of that language. The speaker's possession of such information is assumed, even on the most primitive associationistic views of verbal behavior, for according to such views the speaker's ability to produce coherent utterances depends upon his exploitation of knowledge about transition probabilities between elements of his language.

A second point that cannot be seriously questioned is that some, at least, of the speaker's information about his language must be learned. It is a platitude that no one is born talking, and it seems self-evident that a period of learning is a precondition of fluency, whatever maturational processes may also be involved. No child talks without having been exposed to the utterances of fluent speakers, and the language and dialect he eventually speaks are precisely the language and dialect

to which he is exposed. It should be noticed, however, that to say that the child must learn his language is not to say that the child must be conditioned to speak his language. What kinds of learning mechanisms may be supposed to be involved in mastering the syntax of a first language is the question to which this paper is addressed.

A third point that ought to be treated as self-evident (though it often is not) is that the child must bring to the language-learning situation some amount of intrinsic structure. This structure may take the form of general learning principles or it may take the form of relatively detailed and language-specific information about the kind of grammatical system that underlies natural languages. But what cannot be denied is that any organism that extrapolates from its experience does so on the basis of principles that are not themselves supplied by its experience.

It will be noticed that in talking about the child's contribution to language-learning, I referred to "intrinsic" rather than to "innate" structure. This is because it is at least conceivable that the learning principles in terms of which the child organizes his linguistic experience are themselves learned; for example, that the child is born with a very general capacity to learn learning principles and that it is such learned principles that the child brings to the problem of mastering his language. I do not, in fact, think that suggestion is true in any significant sense, but that is beside the point. What is important is that the task of characterizing the information the child brings to language-learning is, at least in principle, distinguishable from the question of whether the child's intrinsic information is innate. In this paper I shall be concerned solely with the former problem.

How, then, can we proceed in investigating the psychological processes involved in the assimilation of the syntax of a first language? Clearly there are three terms to the relation under consideration. In the first place, there is a body of *observations* that the child must be assumed to make, a body of data about his language provided by the child's exposure to the verbalizations of adults, siblings, and so on. Second, there are whatever *learning principles* the child employs to organize and extrapolate these observations. Third, there is the body of *linguistic information* provided by the application of the principles to the data, the body of knowledge about the structure of his language that the child-cum-fluent-speaker will employ in speaking and understanding the language.

What we say about any of these must condition what we say about each of the others. For example, the child's data plus his intrinsic structure must jointly determine the linguistic information at which he arrives. Hence it is a conclusive disproof of any theory about the child's intrinsic structure to demonstrate that a device having that structure could not learn the syntax of a language on the basis of the kinds of data that the child's verbal environment provides. Suppose, for example, that someone were to maintain that all that is required for language-learning is that the child should be able and inclined to imitate the utterances he hears. It would be a conclusive refutation of that suggestion to point out that a speaker's information about his language is sufficiently rich to permit him to produce or understand novel utterances *ad libitum*. For the fact that the child will learn to deal with utterances of a type not found in his data entails that something must be involved in language-learning other than the imitation of the utterances that are found there.

In short, a comparison of the child's data with a formulation of the linguistic information necessary to speak the language the child learns permits us to estimate the nature and complexity of the child's intrinsic structure. If the information in the child's data closely approximates the linguistic information he must master, we may assume that the role of intrinsic structure is relatively insignificant. Conversely, if the linguistic information at which the child arrives is only indirectly and abstractly related to the data provided by the child's exposure to adult speech, we shall have to suppose that the child's intrinsic structure is correspondingly complex. We have already seen that the two limiting theories can be dismissed with some confidence. On the one hand, it is inconceivable that the child's data contribute no linguistic information, for this would mean that all such information is intrinsic. On the other hand, it seems that the data cannot contribute all the relevant information, for this would be logically incompatible with the fact that the child eventually learns to deal with utterances of sentences he has not previously encountered. It appears that the theory we want must lie somewhere between the two.

To summarize: I propose to consider the child to be a "black box" that converts some body of data about a language into whatever syntactic information is required to speak the language. By comparing what is known about the input to this device with what is known

ut its output, something about its manner of operation and internalization may perhaps be deduced.

#### *E ANALYSIS OF LANGUAGE INPUT*

What is known about these inputs and outputs? About certainures of the latter a good deal can be said as a result of recentk in linguistics. I shall return to this presently. For the moment, us consider what kinds of data the child's encounters with fluent speakers may be assumed to provide.

In the first place, the child gets a *corpus*. That is, he gets a sample of the kinds of utterances fluent speakers of his language typically use. It is conceivable that this sample is biased in certain respects on comparison to a purely random sample. Thus, it has been suggested that speakers addressing children often consciously simplify their utterances both in point of vocabulary choice and in point of syntactic structure; and it is quite certain that adult speakers often complicate induction of the morphology the child must learn by indulging in *r* talk. Research is now being carried out in an attempt to determine precise character of the verbal environment of the child. Until results of this research are known, however, it would be methodologically sound to assume that the child's increasing linguistic proficiency is not to be attributed to any significant extent to the special character of the utterances he hears. For if it is true that utterances are typically directed to children tend to be syntactically simple and children now and then receive the benefit of language tuition in form of corrections of their incorrect utterances, it is equally true much of what children hear is overheard and that all normal children learn to speak, though the differences in the amount of special attention and conscious tuition children receive must vary enormously with variables like social class and birth order.

One point about the corpus should, however, be noticed. If it is something like a randomly selected corpus of adult utterances, it must contain a very substantial number of false starts, slips, grammatical mistakes, and so forth. Most of these the adult speaker is capable of recognizing as distortions of his dialect, hence the attainment of this city is part of what the child must master in learning to speak dialect. Thus the description of the child's task as that of extrapolating from the utterances in his corpus to the sentences of his

language makes the task seem simpler than it is. The child's problem is rather to determine which proper subset of the utterances he hears constitute utterances of sentences and to extrapolate that subset. To put it slightly differently, we may think of the linguistic information at which the child arrives as constituting *inter alia* a theory of the regularities in its corpus, a theory about which such regularities are of systematic significance and which are "accidental." The optimal theory need not count every utterance in the corpus as regular because, as a matter of fact, much of what the child hears is almost certain to violate one or another of the grammatical rules that define the dialect from which the corpus is drawn.

The similarities between the child's problem and normal problems of scientific induction are thus very striking. Like the scientist, the child finds himself with a finite body of observations, some of which are almost certain to be unsystematic. His problem is to discover regularities in these data that, at very least, can be relied upon to hold however much additional data is added. Characteristically the extrapolation takes the form of the construction of a theory that simultaneously marks the systematic similarities among the data at various levels of abstraction, permits the rejection of some of the observational data as unsystematic, and automatically provides a general characterization of the possible future observations. In the case of the learning of language, this theory is precisely the linguistic information at which the child arrives by applying his intrinsic information to the analysis of the corpus. In particular, this linguistic information is at very least required to provide an abstract account of syntactic structure in terms of which systematically relevant features of the observed utterances can be marked, in terms of which some of the observations can be discarded as violating the formation rules of the dialect, and in terms of which the notion "possible sentence of the language" can be defined.

It must be noted that the child's linguistic environment provides him with more than a corpus of utterances. It also provides him with correlations that obtain among members of the corpus and between members of the corpus and various nonlinguistic events. On the one hand, most discourses clearly possess structure beyond the sentence level *Good morning. Isn't it a beautiful day?* but probably not *Good morning. Isn't it a beautiful evening?* On the other hand, many of the assertions the child hears must be true, many of the things he hears

ferred to must exist, many of the questions he hears asked must be answerable, and many of the commands he receives must be performable. Clearly the child could not learn to talk if adults talked at random.

That the child takes advantage of such correlations in learning the semantic system of his language can hardly be doubted. It is also quite conceivable that they may provide him with clues for the learning of one or another syntactic relation. To say that the syntax of a language can be *represented* independently of a representation of its semantics is by no means to claim that the systems are *learned* independently. The difficulty with relying upon "semantic" considerations in explaining language-learning is not, then, that such considerations are known to be irrelevant but simply that we do not know how to describe them in any very revealing way. Which of the indefinitely many correlations between features of language and features of "the world" are relevant to a systematic account of language-learning (or, or that matter, to a systematic account of meaning) is a currently unanswered question. What is perfectly clear, however, is that the theory is enormously more complicated than has usually been realized by psycholinguists. If it be said that the learning-theoretic accounts of reference psychologists have proposed have only been intended as a first step, it must be replied that they are quite certainly a first step in the wrong direction (cf. Chomsky, 1959; Fodor, 1965).

Fortunately, the question of how the child exploits correlations between features of the utterances in his corpus and features of the linguistic environment is probably irrelevant to the problem with which this paper is mainly concerned. The sort of syntactic information that provides the deepest problems for the theory of language-learning is precisely that which concerns the abstract formal structures underlying sentences rather than obviously and immediately meaningful units, such as morphemes and words. To understand this, let us turn from a discussion of the data that provide an input to the child's intrinsic language-learning principles to a discussion of the theory of is language that constitutes their output.

Consider the following linguistically trivial example. For each English active sentence with a certain sort of transitive verb (including *eat*, *ite*, *catch*, . . . etc. but excluding *cost*, *weigh*, . . . etc.), there exists a corresponding passive sentence with the same verb but with the subject and object interchanged. Hence, corresponding to *the dog eats*

*the meat*, there is *the meat was eaten by the dog*; corresponding to *the wolf eats the serfs*, there is *the serfs were eaten by the wolf*, and so on.

It is obvious that the linguistic information at which the child will arrive is sufficient to specify both the indefinite set of actives and the indefinite set of corresponding passives. Nor does it seem plausible to suppose that, having learned that *eat* can take *dog* as subject and *meat* as object in the active, the child must then learn as an entirely independent fact that *eat* can take *dog* as object and *meat* as subject in the passive. On the contrary, an adult speaker, given a novel active, can supply the corresponding passive without further information. It thus seems undeniable that the information that allows the speaker to construct the one form is intimately related to the information that allows him to construct the other.

A reasonable first guess would be that the child learns rules for constructing actives and further rules for converting them into their corresponding passives. For example, having constructed *the dog eats the meat*, the child knows that to form the passive he must interchange the subject and object, attach *en by* to the main verb, and introduce the appropriate inflection of *be* in the auxiliary position. Indeed, most psychologists who have discussed transformations seem to have supposed that some such rule for converting actives into passives is a paradigm for the syntactic characterization of intersentential relations.

That this is a misunderstanding can be seen from the second example just mentioned. Corresponding to *the wolf ate the serfs*, we have *the serfs were eaten by the wolf* not *the serfs was eaten by the wolf*. But notice that it is this latter form that would be produced by a rule which derived the passive from its corresponding active. The difficulty is that in both the active and the passive the number of the verb agrees with the number of its subject. Hence we have the wrong agreement if we allow the passive to be derived from the active in cases where the subject and object differ in number.

From a linguistic point of view, the solution of this problem is simple enough. We hypothesize an underlying base structure in which all the components of the sentence are represented and appropriately marked (*subject*, *verb*, *object*, and so on) but in which the verb is not inflected for number. That is, we postulate a base form that is not a sentence, hence a *fortiori* neither an active nor a passive sentence. It is this underlying form to which the transformational operations are assumed to apply. In particular, the underlying form is converted into the

ve by *inter alia* supplying the appropriate number inflection for the ) and into the passive by first interchanging the subject and object then applying the rule that inflects the verb for number. Since the rangement of subject and object occurs prior to inflection, the same : that inflects the verb in the active can also be employed in the sive; namely, the number of the verb agrees with the number of subject.

n short, we conclude that the linguistic information available to the aker must include rules that permit him to construct nonsentential structures from which such related sentence forms as actives and sives (and, of course, questions, imperatives, and so on) are derived. : while the postulation of such base structures provides no partic- : problem for the linguist, it raises profound difficulties for any : ry of language-learning. In particular the question immediately : sents itself: how does the child learn what the correct base structure : a type of sentence is?

Notice that imitation and reinforcement, the two concepts with which ertican psychologists have traditionally approached problems about yuage-learning, are simply useless here. On the one hand, imitation be relevant as a learning mechanism only where the environment he organism provides it with a model of the behaviors it is required earn. But, by definition, the base structures of a language are not mselves possible utterances in the language. On the contrary, the : ruction of a base structure is presumed to be no more than an : rmediate step in the integration of verbal behavior, and it is : refore the verbal behavior, and not the base forms that are assumed : underlie it, that is available to the child as a possible model for : tation.

imilarly, differential reinforcement, insofar as it may be supposed : play some role in learning, must be contingent either upon the : ntaneous occurrence of the desired behavior as part of the organism's : rant repertoire or upon the occurrence of some related operant that : be "shaped" into an approximation of the behavior desired. Since, : ver, base forms are not uttered by children either in operant : bbing or at any other stage of verbalization, the desired behavior is : available for selective reinforcement. Hence, in the case of the : rning of the base structures of a language, the essential prerequisite : operant conditioning is not satisfied.

n short, a problem that is central to understanding the learning of

yntax is that of arriving at a theory of how the child determines the appropriate base structures for the types of sentences that appear in his corpus. However, the peculiarly abstract relation between base structures and sentences unfix any of the usual learning mechanisms for explaining their assimilation.

### THE INDUCTION OF BASE STRUCTURES

The problem of characterizing a device that, given as input a sample of utterances drawn from a natural language, supplies as output a system of rules that correctly assigns base structures to the sentences of the language is considerably beyond the capacities of current psycholinguistic theory. There is, however, an intermediate problem that can and should be investigated.

Given any corpus of utterances, there are indefinitely many possible abstract descriptions it satisfies. Hence there are indefinitely many infinite sets of which the corpus may represent a finite subset and indefinitely many ways of extrapolating from the corpus to some logically possible language. Of these, all but a very small finite number may be presumed to be "phony" because all but a very small finite number of the possible extrapolations from a corpus exploit properties of utterances that are not of the type that enter into systematic linguistic descriptions. For example, such extrapolations might be defined over intensity of the utterances, or over the number of words they contain, or over the height of the person who speaks them, or with respect to the distance in statute miles between the speaker and the Tower of London. Clearly, whatever intrinsic structure the child brings to the language-learning situation must at least be sufficient to preclude the necessity of running through indefinitely many such absurd hypotheses, or the child would stand no chance whatever of arriving at a reasonable representation of the syntax of his language in any reasonably short period of time.

We may therefore think of the problem of characterizing the child's intrinsic information as dividing into two distinguishable stages. The first task is to provide a characterization of a device that is at least guaranteed only to attempt to describe its input in terms of the kinds of relations that are known to be relevant to systematic linguistic description. Such a device would not, in the normal case, be expected to arrive at a unique best description of the syntax of the language from

his corpus is derived or a unique best mapping of the corpus onto set of sentences of that language. Hence a secondary research goal is to be the characterization of a simplicity metric that, given the corpus candidate descriptions such a device would produce, would use among them in terms of such considerations as their formal complexity, their coherence with putative linguistic universals, the with which they can be integrated with semantic descriptions, so forth. The problem of characterizing this simplicity metric is very much harder than the problem of characterizing a device that judges only non-"phony" extrapolations of corpora. Fortunately, ever, it is also much less pressing. The currently urgent problem is to devise a reasonable account of the induction of underlying structure, of mental operation about which traditional learning mechanisms appear to have nothing whatever to say. We can afford to delay the question of how, given a number of competitive inductions all of which operate in terms of systematically significant linguistic relations, we might go about choosing the best one.

Moreover, it is not entirely unreasonable to suggest that these two sets of investigation correspond to the way that children work. It is likely, on the basis of recent investigations of children's speech, that a child runs through a series of candidate syntaxes, all but the of which are discarded because they fail to capture (or because of misdescribe) relevant regularities in his corpus. That is, the child is to revise or abandon such candidate syntaxes is presumably on the basis of their empirical adequacy *vis-à-vis* the increasingly large corpus available to the child and their simplicity, both in respect to the number of rules they require the child to store and in respect to their coherence with nonsyntactic portions of the grammar.

I have nothing to say about the character of the simplicity metric. I think it is possible to make a few comments about the character of the kind of quasi-discovery procedure that I have suggested ought to be used in immediate research goal. That is, I want to make a few proposals about the construction of a device that will at least ignore the indefinitely many, obviously incorrect, extrapolations of a corpus in favor of those extrapolations that project only features of the type known to be linguistically relevant.

It is worth considering the possibility that the child may bring to language-learning situation a set of rules that takes him from the

recognition of specified formal relations within and among the strings in his data to specific putative characterizations of underlying structures for strings of those types. Such rules would implicitly define the space of hypotheses through which the child must search in order to arrive at the precisely correct syntactic analysis of his corpus. Presumably the rules would have to be so formulated as to assure (1) that the number of possible analyses assigned to a given corpus is fairly small; (2) that the correct analysis (or, at any event, a best analysis) is among these; (3) that the rules project no analysis that describes the corpus in terms of the sorts of phony properties already discussed, but that all the analyses exploit only relations of types that sometimes figure in adequate syntactic theories.

To see how this might work, consider the kinds of relations between base structures and terminal strings that obtain in the case of the English auxiliary construction. That construction permits an indefinite class of sentences, of which A and B are subsets, but excludes an indefinite class of strings, of which \*C is a subset.

A

John is eating lunch	John has eaten lunch
they are playing ball	they have been playing ball
John and Mary are seeing a ghost	John and Mary have seen a ghost

B

etc.

etc.

\*C

John has eating lunch
they have being playing ball
John and Mary have seeing a ghost

etc.

In short, when the auxiliary is developed as some inflection of *be*, the main verb takes *ing*, but when the auxiliary is developed as some inflection of *have*, the main verb takes *en*. This relation is expressed in the syntax by assuming (*a*) that "have + en" and "be + ing" are

stituents in the base structure and (*b*) that the base structures of sentences in A and B have undergone a permutation transformation it effects the appropriate reordering, roughly:

$$(\text{have} + \text{en}) + (\text{be} + \text{ing}) + \text{verb} \Rightarrow \left\{ \begin{array}{l} (\text{have}) + \text{verb} + (\text{en}) \\ (\text{be}) + \text{verb} + (\text{ing}) \end{array} \right\}$$

We may reformulate this situation slightly more abstractly. The sentences in A and B all have the abstract description *IXJ*, where the *m* of the *J*th element depends upon the form of the *I*th element and where *X* represents some intervening string. To arrive at the proper syntactic analysis of strings that satisfy this description, the child must assume that (*IJ*) (*X*) represents the constituent analysis of the base structure and that their transformational history includes a transformation

$$(IJ) (X) \Rightarrow IXJ$$

What I wish to suggest as a working hypothesis is that precisely this information is included in (and typical of) the child's intrinsic structure. That is, that the child comes to the language-learning situation with the instruction that whenever he finds sets of terminals that satisfy the description *IXJ* under the conditions just mentioned, he of the preferred hypotheses about their syntactic analysis is that their base structure contains (*IJ*) and that their transformational history contains the operation just cited.

It is worth noticing that this rule also provides for the correct analysis of the class of sentences instanced by such examples as *John phoned Mary up*; *John phoned the girl he knew in Chicago up*; *a friend of mine looked over the car*; *a friend of mine looked the car over*. In these cases the interdependent items *IJ* are a verb with its attached particle, and the *X* is a noun phrase (NP). The usual syntactic treatment of such cases is to assume that the underlying structure is (Verb + Particle) (NP) and that the terminal sequence produced by the application of an (optional) transformation that mutates the particle with the NP.

There are three major points I want to make about the suggestion implicit in these examples. The first is that the child would clearly be able to employ the sort of information I am supposing to be intrinsic unless he was able to recognize a relation of abstract conformity between elements and/or sequences of elements in his corpus. In the

present cases, employing the relevant rule for inducing base structures requires his having available the morphological information that *has* and *have* on the one hand and *is* and *are* on the other are both representations of the same underlying elements, *have* in the former case and *be* in the latter. Analogously, he could not employ the relevant rule for analyzing constructions like *John phoned Mary up* unless he had available the syntactic information that *phone up*, *look over*, and so on all have the abstract representation "Verb Particle." Without such information, it would be impossible for the child to analyze the sentences in his corpus as *IXJ* and hence impossible for him to apply the sort of rule suggested above for arriving at their base structures.

It may be that the techniques of substitution and classification traditionally employed in attempts to formulate linguistic discovery procedure will prove useful here. It is, of course, notorious that such procedures do not arrive at unique correct grammars for arbitrary finite corpora. But the present point is that they need not do so in order to aid the child in arriving at tentative hypotheses about conformity relations. I am proposing, rather, that the child may employ such relations as substitutability-in-frames to arrive at tentative classifications of elements and sequences of elements in his corpus and hence at tentative domains for the application of intrinsic rules for inducing base structures. Whether a given such classification is retained or discarded would be contingent upon the relative simplicity of the entire system of which it forms a part. Thus, if the child arrives at a taxonomy of a set of strings such that each string in the set has the representation *IXJ*, then if he has some independent reason for supposing that the base form for these strings is not (*IJ*) (*X*), one of the options available to him is to abandon the original taxonomy. He might, for example, suppose that what appear to be recurrences of conforming *I* elements are in fact occurrences of syntactically unrelated homonyms.

In short, I am suggesting a process for the induction of base structures that includes the following phases: (1) provide tentative abstract representations of the derived structure of the corpus by employing techniques that would, presumably, depend in part upon the assumption that distributionally similar sequences often belong to the same class; (2) provide tentative base structures for each such representation by employing intrinsic rules like the one already



putative analysis of a certain set of strings in the data is a function of the relative simplicity of the entire syntactic theory of the language the child eventually arrives at. What makes the analysis of strings in C as permuted forms of *IXI* unacceptable is simply that it would require that, in the generation of the strings in C, we use a number agreement rule that differs from the number agreement rule employed in generating strings in D, and this proliferation of rules is in conflict with general principles of parsimony. (The example also suggests, by the way, that if parents do simplify the syntax of their speech when they address children, they may thereby make it harder for the child to learn the correct syntactic analysis of his language. Rules that hold for selected sets of simple sentences may have to be abandoned in the light of examples of sentences of more complicated types. The view that one makes language-learning easier by unsystematically limiting the child's access to data about his language is a decidedly peculiar one.)

I do not mind that using the sorts of rules I am supposing may be intrinsic will often lead the child to the wrong syntactic analysis. For, once again, I am not attempting to answer the extremely difficult question, "What sort of device would project a unique correct grammar on the basis of exposure to a corpus?" I am only attempting to formulate a strategy for answering the much easier question, "What sort of device would project candidate grammars that are reasonably sensitive to the contents of the corpus and that operate only with the sorts of relations that are known to figure in linguistic descriptions?" What the example just discussed shows is that if we want the device to handle dependencies in a reasonably sophisticated way, it will be required to take account of more subtle relations among terminals than the ones it needs to notice in projecting underlying structures for constructions of the form "Verb NP Particle." In particular, it will have to take account of the difference between a corpus in which the strings can be analyzed *IXI* and a corpus in which the strings can be analyzed *IXI<sub>1</sub>, IXI<sub>2</sub>, . . . , IXI<sub>n</sub>* and employ correspondingly different inference rules for the induction of the underlying forms of these two types of sets.

The final point I should like to make about this proposal is that it is surely incorrect as stated. What seems to me reasonable as a research strategy is to attempt to analyze the kinds of relations well-confirmed syntaxes claim obtain between strings of specified formal character and their respective underlying structures, and to work on

the assumption that the intrinsic structure of children includes representations of such relations. Clearly this view is capable in principle of accounting for the child's induction of syntactic analyses far more complicated than those he could arrive at solely by applying such restricted techniques as substitution and classification. It is thus a virtue of the proposal that it makes the induction of underlying structure appear a little less mysterious than it has sometimes seemed and thus takes the sting out of the objection that transformational grammars could not be psychologically revealing because they hypothesize relations too complicated for a child to learn.

Moreover, even the linguistically trivial examples already mentioned suggest that an appropriate formulation of the rules that map from features of the strings in the corpus to features of their underlying structure might permit the homogeneous (and hence parsimonious) treatment of a fairly wide variety of superficially heterogeneous structures. Thus, the rule I mentioned would account for both the child's ability to induce the underlying structure of "Verb Particle" constructions and his ability to handle certain agreements in the auxiliary. For, though constructions like *have eaten* are syntactically unrelated to constructions like *phone him up*, both such constructions bear the same abstract relation to their respective base forms, and I have assumed that the child's intrinsic information permits him to take account of that fact.

But if such examples reveal the virtues of the present proposal, they also reveal its difficulties. What one would need to make it interesting would be a more general characterization of the kinds of relations that obtain between base forms and derived forms than is now available. To make the proposal plausible one would need to show that a reasonably wide variety of such relations can be induced on the basis of a relatively small amount of intrinsic information of an abstract and general sort. The difference between the present proposal and the *ad hoc* suggestion that the child is born with a list of the solutions to all the language-learning problems he will encounter resides, after all, precisely in the generality and simplicity with which the presumed intrinsic information can be specified. At the present stage I am unable to do more than provide examples of the kinds of inference rules that might work for certain specific kinds of cases. I am sure the examples must be wrong as stated, and I do not know how to state them so that they are at once more powerful and more precise. But I suspect

that that would be a good problem to think about when one is inclined to speculate on the role of intrinsic structure in first-language learning.

## REFERENCES

- CHOMSKY, N. A review of Skinner's *Verbal Behavior*. *Language*, 1959, 35, 26.  
 FODOR, J. A. Could meaning be an  $r_m$ ? *J. verb. Learn. and verb. Behav.*, 1965, 4, 2.

## GENERAL DISCUSSION

### Fodor presentation

1. An insistent question concerned the difference, if any, between Fodor's proposal and that of McNeill. Several participants, including McNeill, felt that Fodor's position would reduce to McNeill's. Fodor acknowledged that the distinction might well be between two senses of parsimony. McNeill's child had a lot of structural information wired in, for example about the kinds of trees (patterns of hierarchical organization) relating grammatical classes, and Fodor's child was equipped with a set of inference rules permitting the generation of a variety of candidate grammars, all of which must share formal universals like the possession of a hierarchically organized component. This led McNeill to wonder whether Fodor was in fact attributing less foreknowledge to the child than he was; Fodor replied that his child was equipped with no structures at all, for example, no characterizations of the abstract form of particular trees. Rather, the child was assumed to possess rules for processing the information in his corpus such that the consequence of applying those rules was the generation of trees of the relevant shape. This might turn out to be the same proposal as McNeill's, but superficially it did not seem to be, for although Fodor's proposal did not imply less prestructure than McNeill's, it implied that the innate information was available in a rather different form from the one which McNeill proposed. Given the right set of inference rules, the problem of the formal facts of language might be resolved automatically.

2. Another persistent discussion topic centered on the fact that in

# THE GENESIS OF LANGUAGE

## A Psycholinguistic Approach

*Proceedings of a conference on "Language Development in Children,"  
sponsored by the National Institute of Child Health and Human  
Development, National Institutes of Health*

Edited by Frank Smith and George A. Miller



THE M.I.T. PRESS  
Massachusetts Institute of Technology  
Cambridge, Massachusetts, and London, England