

Methodological Preliminaries

§ 1. GENERATIVE GRAMMARS AS THEORIES OF LINGUISTIC COMPETENCE

THIS study will touch on a variety of topics in syntactic theory and English syntax, a few in some detail, several quite superficially, and none exhaustively. It will be concerned with the syntactic component of a generative grammar, that is, with the rules that specify the well-formed strings of minimal syntactically functioning units (*formatives*) and assign structural information of various kinds both to these strings and to strings that deviate from well-formedness in certain respects.

The general framework within which this investigation will proceed has been presented in many places, and some familiarity with the theoretical and descriptive studies listed in the bibliography is presupposed. In this chapter, I shall survey briefly some of the main background assumptions, making no serious attempt here to justify them but only to sketch them clearly.

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance. This seems to me to have been the position of the founders of modern general linguistics, and no cogent reason for

modifying it has been offered. To study actual linguistic performance, we must consider the interaction of a variety of factors, of which the underlying competence of the speaker-hearer is only one. In this respect, study of language is no different from empirical investigation of other complex phenomena.

We thus make a fundamental distinction between *competence* (the speaker-hearer's knowledge of his language) and *performance* (the actual use of language in concrete situations). Only under the idealization set forth in the preceding paragraph is performance a direct reflection of competence. In actual fact, it obviously could not directly reflect competence. A record of natural speech will show numerous false starts, deviations from rules, changes of plan in mid-course, and so on. The problem for the linguist, as well as for the child learning the language, is to determine from the data of performance the underlying system of rules that has been mastered by the speaker-hearer and that he puts to use in actual performance. Hence, in the technical sense, linguistic theory is mentalistic, since it is concerned with discovering a mental reality underlying actual behavior.¹ Observed use of language or hypothesized dispositions to respond, habits, and so on, may provide evidence as to the nature of this mental reality, but surely cannot constitute the actual subject matter of linguistics, if this is to be a serious discipline. The distinction I am noting here is related to the *langue-parole* distinction of Saussure; but it is necessary to reject his concept of *langue* as merely a systematic inventory of items and to return rather to the Humboldtian conception of underlying competence as a system of generative processes. For discussion, see Chomsky (1964).

A grammar of a language purports to be a description of the ideal speaker-hearer's intrinsic competence. If the grammar is, furthermore, perfectly explicit — in other words, if it does not rely on the intelligence of the understanding reader but rather provides an explicit analysis of his contribution — we may (somewhat redundantly) call it a *generative grammar*.

A fully adequate grammar must assign to each of an infinite range of sentences a structural description indicating how this

sentence is understood by the ideal speaker-hearer. This is the traditional problem of descriptive linguistics, and traditional grammars give a wealth of information concerning structural descriptions of sentences. However, valuable as they obviously are, traditional grammars are deficient in that they leave unexpressed many of the basic regularities of the language with which they are concerned. This fact is particularly clear on the level of syntax, where no traditional or structuralist grammar goes beyond classification of particular examples to the stage of formulation of generative rules on any significant scale. An analysis of the best existing grammars will quickly reveal that this is a defect of principle, not just a matter of empirical detail or logical preciseness. Nevertheless, it seems obvious that the attempt to explore this largely uncharted territory can most profitably begin with a study of the kind of structural information presented by traditional grammars and the kind of linguistic processes that have been exhibited, however informally, in these grammars.²

The limitations of traditional and structuralist grammars should be clearly appreciated. Although such grammars may contain full and explicit lists of exceptions and irregularities, they provide only examples and hints concerning the regular and productive syntactic processes. Traditional linguistic theory was not unaware of this fact. For example, James Beattie (1788) remarks that

Languages, therefore, resemble men in this respect, that, though each has peculiarities, whereby it is distinguished from every other, yet all have certain qualities in common. The peculiarities of individual tongues are explained in their respective grammars and dictionaries. Those things, that all languages have in common, or that are necessary to every language, are treated of in a science, which some have called *Universal* or *Philosophical* grammar.

Somewhat earlier, Du Marsais defines universal and particular grammar in the following way (1729; quoted in Sahlín, 1928, pp. 29-30):

Il y a dans la grammaire des observations qui conviennent à toutes les langues; ces observations forment ce qu'on appelle la grammaire

générale: telles sont les remarques que l'on a faites sur les sons articulés, sur les lettres qui sont les signes de ces sons; sur la nature des mots, et sur les différentes manières dont ils doivent être ou arrangés ou terminés pour faire un sens. Outre ces observations générales, il y en a qui ne sont propres qu'à une langue particulière; et c'est ce qui forme les grammaires particulières de chaque langue.

Within traditional linguistic theory, furthermore, it was clearly understood that one of the qualities that all languages have in common is their "creative" aspect. Thus an essential property of language is that it provides the means for expressing indefinitely many thoughts and for reacting appropriately in an indefinite range of new situations (for references, cf. Chomsky, 1964, forthcoming). The grammar of a particular language, then, is to be supplemented by a universal grammar that accommodates the creative aspect of language use and expresses the deep-seated regularities which, being universal, are omitted from the grammar itself. Therefore it is quite proper for a grammar to discuss only exceptions and irregularities in any detail. It is only when supplemented by a universal grammar that the grammar of a language provides a full account of the speaker-hearer's competence.

Modern linguistics, however, has not explicitly recognized the necessity for supplementing a "particular grammar" of a language by a universal grammar if it is to achieve descriptive adequacy. It has, in fact, characteristically rejected the study of universal grammar as misguided; and, as noted before, it has not attempted to deal with the creative aspect of language use. It thus suggests no way to overcome the fundamental descriptive inadequacy of structuralist grammars.

Another reason for the failure of traditional grammars, particular or universal, to attempt a precise statement of regular processes of sentence formation and sentence interpretation lay in the widely held belief that there is a "natural order of thoughts" that is mirrored by the order of words. Hence, the rules of sentence formation do not really belong to grammar but to some other subject in which the "order of thoughts" is studied. Thus in the *Grammaire générale et raisonnée* (Lancelot

et al., 1660) it is asserted that, aside from figurative speech, the sequence of words follows an "ordre naturel," which conforms "à l'expression naturelle de nos pensées." Consequently, few grammatical rules need be formulated beyond the rules of ellipsis, inversion, and so on, which determine the figurative use of language. The same view appears in many forms and variants. To mention just one additional example, in an interesting essay devoted largely to the question of how the simultaneous and sequential array of ideas is reflected in the order of words, Diderot concludes that French is unique among languages in the degree to which the order of words corresponds to the natural order of thoughts and ideas (Diderot, 1751). Thus "quel que soit l'ordre des termes dans une langue ancienne ou moderne, l'esprit de l'écrivain a suivi l'ordre didactique de la syntaxe française" (p. 390); "Nous disons les choses en français, comme l'esprit est forcé de les considérer en quelque langue qu'on écrive" (p. 371). With admirable consistency he goes on to conclude that "notre langue *pédestre* a sur les autres l'avantage de l'utile sur l'agréable" (p. 372); thus French is appropriate for the sciences, whereas Greek, Latin, Italian, and English "sont plus avantageuses pour les lettres." Moreover,

le bon sens choisirait la langue française; mais . . . l'imagination et les passions donneront la préférence aux langues anciennes et à celles de nos voisins . . . il faut parler français dans la société et dans les écoles de philosophie; et grec, latin, anglais, dans les chaires et sur les théâtres; . . . notre langue sera celle de la vérité, si jamais elle revient sur la terre; et . . . la grecque, la latine et les autres seront les langues de la fable et du mensonge. Le français est fait pour instruire, éclairer et convaincre; le grec, le latin, l'italien, l'anglais, pour persuader, émouvoir et tromper: parlez grec, latin, italien au peuple; mais parlez français au sage. (pp. 371-372)

In any event, insofar as the order of words is determined by factors independent of language, it is not necessary to describe it in a particular or universal grammar, and we therefore have principled grounds for excluding an explicit formulation of syntactic processes from grammar. It is worth noting that this naïve view of language structure persists to modern times in

various forms, for example, in Saussure's image of a sequence of expressions corresponding to an amorphous sequence of concepts or in the common characterization of language use as merely a matter of use of words and phrases (for example, Ryle, 1953).

But the fundamental reason for this inadequacy of traditional grammars is a more technical one. Although it was well understood that linguistic processes are in some sense "creative," the technical devices for expressing a system of recursive processes were simply not available until much more recently. In fact, a real understanding of how a language can (in Humboldt's words) "make infinite use of finite means" has developed only within the last thirty years, in the course of studies in the foundations of mathematics. Now that these insights are readily available it is possible to return to the problems that were raised, but not solved, in traditional linguistic theory, and to attempt an explicit formulation of the "creative" processes of language. There is, in short, no longer a technical barrier to the full-scale study of generative grammars.

Returning to the main theme, by a generative grammar I mean simply a system of rules that in some explicit and well-defined way assigns structural descriptions to sentences. Obviously, every speaker of a language has mastered and internalized a generative grammar that expresses his knowledge of his language. This is not to say that he is aware of the rules of the grammar or even that he can become aware of them, or that his statements about his intuitive knowledge of the language are necessarily accurate. Any interesting generative grammar will be dealing, for the most part, with mental processes that are far beyond the level of actual or even potential consciousness; furthermore, it is quite apparent that a speaker's reports and viewpoints about his behavior and his competence may be in error. Thus a generative grammar attempts to specify what the speaker actually knows, not what he may report about his knowledge. Similarly, a theory of visual perception would attempt to account for what a person actually sees and the mechanisms that determine this rather than his statements about what he sees and why, though these state-

ments may provide useful, in fact, compelling evidence for such a theory.

To avoid what has been a continuing misunderstanding, it is perhaps worth while to reiterate that a generative grammar is not a model for a speaker or a hearer. It attempts to characterize in the most neutral possible terms the knowledge of the language that provides the basis for actual use of language by a speaker-hearer. When we speak of a grammar as generating a sentence with a certain structural description, we mean simply that the grammar assigns this structural description to the sentence. When we say that a sentence has a certain derivation with respect to a particular generative grammar, we say nothing about how the speaker or hearer might proceed, in some practical or efficient way, to construct such a derivation. These questions belong to the theory of language use—the theory of performance. No doubt, a reasonable model of language use will incorporate, as a basic component, the generative grammar that expresses the speaker-hearer's knowledge of the language; but this generative grammar does not, in itself, prescribe the character or functioning of a perceptual model or a model of speech production. For various attempts to clarify this point, see Chomsky (1957), Gleason (1961), Miller and Chomsky (1963), and many other publications.

Confusion over this matter has been sufficiently persistent to suggest that a terminological change might be in order. Nevertheless, I think that the term "generative grammar" is completely appropriate, and have therefore continued to use it. The term "generate" is familiar in the sense intended here in logic, particularly in Post's theory of combinatorial systems. Furthermore, "generate" seems to be the most appropriate translation for Humboldt's term *erzeugen*, which he frequently uses, it seems, in essentially the sense here intended. Since this use of the term "generate" is well established both in logic and in the tradition of linguistic theory, I can see no reason for a revision of terminology.

§ 2. TOWARD A THEORY OF PERFORMANCE

There seems to be little reason to question the traditional view that investigation of performance will proceed only so far as understanding of underlying competence permits. Furthermore, recent work on performance seems to give new support to this assumption. To my knowledge, the only concrete results that have been achieved and the only clear suggestions that have been put forth concerning the theory of performance, outside of phonetics, have come from studies of performance models that incorporate generative grammars of specific kinds — that is, from studies that have been based on assumptions about underlying competence.³ In particular, there are some suggestive observations concerning limitations on performance imposed by organization of memory and bounds on memory, and concerning the exploitation of grammatical devices to form deviant sentences of various types. The latter question is one to which we shall return in Chapters 2 and 4. To clarify further the distinction between competence and performance, it may be useful to summarize briefly some of the suggestions and results that have appeared in the last few years in the study of performance models with limitations of memory, time, and access.

For the purposes of this discussion, let us use the term "acceptable" to refer to utterances that are perfectly natural and immediately comprehensible without paper-and-pencil analysis, and in no way bizarre or outlandish. Obviously, acceptability will be a matter of degree, along various dimensions. One could go on to propose various operational tests to specify the notion more precisely (for example, rapidity, correctness, and uniformity of recall and recognition, normalcy of intonation).⁴ For present purposes, it is unnecessary to delimit it more carefully. To illustrate, the sentences of (1) are somewhat more acceptable, in the intended sense, than those of (2):

- (1) (i) I called up the man who wrote the book that you told me about
 (ii) quite a few of the students who you met who come from New York are friends of mine

§ 2. TOWARD A THEORY OF PERFORMANCE

- (iii) John, Bill, Tom, and several of their friends visited us last night
 (2) (i) I called the man who wrote the book that you told me about up
 (ii) the man who the boy who the students recognized pointed out is a friend of mine

The more acceptable sentences are those that are more likely to be produced, more easily understood, less clumsy, and in some sense more natural.⁵ The unacceptable sentences one would tend to avoid and replace by more acceptable variants, wherever possible, in actual discourse.

The notion "acceptable" is not to be confused with "grammatical." Acceptability is a concept that belongs to the study of performance, whereas grammaticalness belongs to the study of competence. The sentences of (2) are low on the scale of acceptability but high on the scale of grammaticalness, in the technical sense of this term. That is, the generative rules of the language assign an interpretation to them in exactly the way in which they assign an interpretation to the somewhat more acceptable sentences of (1). Like acceptability, grammaticalness is, no doubt, a matter of degree (cf. Chomsky, 1955, 1957, 1961), but the scales of grammaticalness and acceptability do not coincide. Grammaticalness is only one of many factors that interact to determine acceptability. Correspondingly, although one might propose various operational tests for acceptability, it is unlikely that a necessary and sufficient operational criterion might be invented for the much more abstract and far more important notion of grammaticalness. The unacceptable grammatical sentences often cannot be used, for reasons having to do, not with grammar, but rather with memory limitations, intonational and stylistic factors, "iconic" elements of discourse (for example, a tendency to place logical subject and object early rather than late; cf. note 32, Chapter 2, and note 9, Chapter 3), and so on. Note that it would be quite impossible to characterize the unacceptable sentences in grammatical terms. For example, we cannot formulate particular rules of the grammar in such a way as

to exclude them. Nor, obviously, can we exclude them by limiting the number of reapplications of grammatical rules in the generation of a sentence, since unacceptability can just as well arise from application of distinct rules, each being applied only once. In fact, it is clear that we can characterize unacceptable sentences only in terms of some "global" property of derivations and the structures they define — a property that is attributable, not to a particular rule, but rather to the way in which the rules interrelate in a derivation.

This observation suggests that the study of performance could profitably begin with an investigation of the acceptability of the simplest formal structures in grammatical sentences. The most obvious formal property of utterances is their bracketing into constituents of various types, that is, the "tree structure" associated with them. Among such structures we can distinguish various kinds — for example, those to which we give the following conventional technical names, for the purposes of this discussion:

- (3) (i) nested constructions
- (ii) self-embedded constructions
- (iii) multiple-branching constructions
- (iv) left-branching constructions
- (v) right-branching constructions

The phrases *A* and *B* form a nested construction if *A* falls totally within *B*, with some nonnull element to its left within *B* and some nonnull element to its right within *B*. Thus the phrase "the man who wrote the book that you told me about" is nested in the phrase "called the man who wrote the book that you told me about up," in (2i). The phrase *A* is self-embedded in *B* if *A* is nested in *B* and, furthermore, *A* is a phrase of the same type as *B*. Thus "who the students recognized" is self-embedded in "who the boy who the students recognized pointed out," in (2ii), since both are relative clauses. Thus nesting has to do with bracketing, and self-embedding with labeling of brackets as well. A multiple-branching construction is one with no internal structure. In (1iii), the Subject Noun Phrase is multiple-branch-

ing, since "John," "Bill," "Tom," and "several of their friends" are its immediate constituents, and have no further association among themselves. In terms of bracketing, a multiple-branching construction has the form $[[A][B]\cdots[M]]$. A left-branching structure is of the form $[[[\cdots]\cdots]\cdots]$ — for example, in English, such indefinitely iterable structures as $[[[[[John]'s\ brother]'s\ father]'s\ uncle]$ or $[[[the\ man\ who\ you\ met]\ from\ Boston]\ who\ was\ on\ the\ train]$, or (1ii), which combines several kinds of left-branching. Right-branching structures are those with the opposite property — for example, the Direct-Object of (1i) or $[this\ is\ [the\ cat\ that\ caught\ [the\ rat\ that\ stole\ the\ cheese]]]$.

The effect of these superficial aspects of sentence structure on performance has been a topic of study since almost the very inception of recent work on generative grammar, and there are some suggestive observations concerning their role in determining acceptability (that is, their role in limiting performance). Summarizing this work briefly, the following observations seem plausible:

- (4) (i) repeated nesting contributes to unacceptability
- (ii) self-embedding contributes still more radically to unacceptability
- (iii) multiple-branching constructions are optimal in acceptability
- (iv) nesting of a long and complex element reduces acceptability
- (v) there are no clear examples of unacceptability involving only left-branching or only right-branching, although these constructions are unnatural in other ways — thus, for example, in reading the right-branching construction "this is the cat that caught the rat that stole the cheese," the intonation breaks are ordinarily inserted in the wrong places (that is, after "cat" and "rat," instead of where the main brackets appear)

In some measure, these phenomena are easily explained. Thus it is known (cf. Chomsky, 1959a; and for discussion, Chomsky, 1961, and Miller and Chomsky, 1963) that an optimal perceptual

device, even with a bounded memory, can accept unbounded left-branching and right-branching structures, though nested (hence ultimately self-embedded) structures go beyond its memory capacity. Thus case (4i) is simply a consequence of finiteness of memory, and the unacceptability of such examples as (2ii) raises no problem.

If (4ii) is correct,⁶ then we have evidence for a conclusion about organization of memory that goes beyond the triviality that it must be finite in size. An optimal finite perceptual device of the type discussed in Chomsky (1959a) need have no more difficulty with self-embedding than with other kinds of nesting (see Bar-Hillel, Kasher, and Shamir, 1963, for a discussion of this point). To account for the greater unacceptability of self-embedding (assuming this to be a fact), we must add other conditions on the perceptual device beyond mere limitation of memory. We might assume, for example, that the perceptual device has a stock of analytic procedures available to it, one corresponding to each kind of phrase, and that it is organized in such a way that it is unable (or finds it difficult) to utilize a procedure φ while it is in the course of executing φ . This is not a necessary feature of a perceptual model, but it is a rather plausible one, and it would account for (4ii). See, in this connection, Miller and Isard (1964).

The high acceptability of multiple-branching, as in case (4iii), is easily explained on the rather plausible assumption that the ratio of number of phrases to number of formatives (the node-to-terminal node ratio, in a tree-diagram of a sentence) is a rough measure of the amount of computation that has to be performed in analysis. Thus multiple coordination would be the simplest kind of construction for an analytic device — it would impose the least strain on memory.⁷ For discussion, see Miller and Chomsky (1963).

Case (4iv) suggests decay of memory, perhaps, but raises unsolved problems (see Chomsky, 1961, note 19).

Case (4v) follows from the result about optimal perceptual models mentioned earlier. But it is unclear why left- and right-branching structures should become unnatural after a certain point, if they actually do.⁸

One might ask whether attention to less superficial aspects of grammatical structure than those of (3) could lead to somewhat deeper conclusions about performance models. This seems entirely possible. For example, in Miller and Chomsky (1963) some syntactic and perceptual considerations are adduced in support of a suggestion (which is, to be sure, highly speculative) as to the somewhat more detailed organization of a perceptual device. In general, it seems that the study of performance models incorporating generative grammars may be a fruitful study; furthermore, it is difficult to imagine any other basis on which a theory of performance might develop.

There has been a fair amount of criticism of work in generative grammar on the grounds that it slights study of performance in favor of study of underlying competence. The facts, however, seem to be that the only studies of performance, outside of phonetics (but see note 3), are those carried out as a by-product of work in generative grammar. In particular, the study of memory limitations just summarized and the study of deviation from rules, as a stylistic device, to which we return in Chapters 2 and 4, have developed in this way. Furthermore, it seems that these lines of investigation can provide some insight into performance. Consequently, this criticism is unwarranted, and, furthermore, completely misdirected. It is the descriptivist limitation-in-principle to classification and organization of data, to "extracting patterns" from a corpus of observed speech, to describing "speech habits" or "habit structures," insofar as these may exist, etc., that precludes the development of a theory of actual performance.

§ 3. THE ORGANIZATION OF A GENERATIVE GRAMMAR

Returning now to the question of competence and the generative grammars that purport to describe it, we stress again that knowledge of a language involves the implicit ability to understand indefinitely many sentences.⁹ Hence, a generative grammar must be a system of rules that can iterate to generate an in-

definitely large number of structures. This system of rules can be analyzed into the three major components of a generative grammar: the syntactic, phonological, and semantic components.¹⁰

The syntactic component specifies an infinite set of abstract formal objects, each of which incorporates all information relevant to a single interpretation of a particular sentence.¹¹ Since I shall be concerned here only with the syntactic component, I shall use the term "sentence" to refer to strings of formatives rather than to strings of phones. It will be recalled that a string of formatives specifies a string of phones uniquely (up to free variation), but not conversely.

The phonological component of a grammar determines the phonetic form of a sentence generated by the syntactic rules. That is, it relates a structure generated by the syntactic component to a phonetically represented signal. The semantic component determines the semantic interpretation of a sentence. That is, it relates a structure generated by the syntactic component to a certain semantic representation. Both the phonological and semantic components are therefore purely interpretive. Each utilizes information provided by the syntactic component concerning formatives, their inherent properties, and their interrelations in a given sentence. Consequently, the syntactic component of a grammar must specify, for each sentence, a *deep structure* that determines its semantic interpretation and a *surface structure* that determines its phonetic interpretation. The first of these is interpreted by the semantic component; the second, by the phonological component.¹²

It might be supposed that surface structure and deep structure will always be identical. In fact, one might briefly characterize the syntactic theories that have arisen in modern structural (taxonomic) linguistics as based on the assumption that deep and surface structures are actually the same (cf. Postal, 1964a, Chomsky, 1964). The central idea of transformational grammar is that they are, in general, distinct and that the surface structure is determined by repeated application of certain formal operations called "grammatical transformations" to objects of a more

elementary sort. If this is true (as I assume, henceforth), then the syntactic component must generate deep and surface structures, for each sentence, and must interrelate them. This idea has been clarified substantially in recent work, in ways that will be described later. In Chapter 3, I shall present a specific and, in part, new proposal as to precisely how it should be formulated. For the moment, it is sufficient to observe that although the Immediate Constituent analysis (labeled bracketing) of an actual string of formatives may be adequate as an account of surface structure, it is certainly not adequate as an account of deep structure. My concern in this book is primarily with deep structure and, in particular, with the elementary objects of which deep structure is constituted.

To clarify exposition, I shall use the following terminology, with occasional revisions as the discussion proceeds.

The *base* of the syntactic component is a system of rules that generate a highly restricted (perhaps finite) set of *basic strings*, each with an associated structural description called a *base Phrase-marker*. These base Phrase-markers are the elementary units of which deep structures are constituted. I shall assume that no ambiguity is introduced by rules of the base. This assumption seems to me correct, but has no important consequences for what follows here, though it simplifies exposition. Underlying each sentence of the language there is a sequence of base Phrase-markers, each generated by the base of the syntactic component. I shall refer to this sequence as the *basis* of the sentence that it underlies.

In addition to its base, the syntactic component of a generative grammar contains a *transformational* subcomponent. This is concerned with generating a sentence, with its surface structure, from its basis. Some familiarity with the operation and effects of transformational rules is henceforth presupposed.

Since the base generates only a restricted set of base Phrase-markers, most sentences will have a sequence of such objects as an underlying basis. Among the sentences with a single base Phrase-marker as basis, we can delimit a proper subset called "kernel sentences." These are sentences of a particularly simple

sort that involve a minimum of transformational apparatus in their generation. The notion "kernel sentence" has, I think, an important intuitive significance, but since kernel sentences play no distinctive role in generation or interpretation of sentences, I shall say nothing more about them here. One must be careful not to confuse kernel sentences with the basic strings that underlie them. The basic strings and base Phrase-markers do, it seems, play a distinctive and crucial role in language use.

Since transformations will not be considered here in detail, no careful distinction will be made, in the case of a sentence with a single element in its basis, between the basic string underlying this sentence and the sentence itself. In other words, at many points in the exposition I shall make the tacit simplifying (and contrary-to-fact) assumption that the underlying basic string *is* the sentence, in this case, and that the base Phrase-marker is the surface structure as well as the deep structure. I shall try to select examples in such a way as to minimize possible confusion, but the simplifying assumption should be borne in mind throughout.

§ 4. JUSTIFICATION OF GRAMMARS

Before entering directly into an investigation of the syntactic component of a generative grammar, it is important to give some thought to several methodological questions of justification and adequacy.

There is, first of all, the question of how one is to obtain information about the speaker-hearer's competence, about his knowledge of the language. Like most facts of interest and importance, this is neither presented for direct observation nor extractable from data by inductive procedures of any known sort. Clearly, the actual data of linguistic performance will provide much evidence for determining the correctness of hypotheses about underlying linguistic structure, along with introspective reports (by the native speaker, or the linguist who has learned the language). This is the position that is universally adopted in practice, although there are methodological discus-

sions that seem to imply a reluctance to use observed performance or introspective reports as evidence for some underlying reality.

In brief, it is unfortunately the case that no adequate formalizable techniques are known for obtaining reliable information concerning the facts of linguistic structure (nor is this particularly surprising). There are, in other words, very few reliable experimental or data-processing procedures for obtaining significant information concerning the linguistic intuition of the native speaker. It is important to bear in mind that when an operational procedure is proposed, it must be tested for adequacy (exactly as a theory of linguistic intuition — a grammar — must be tested for adequacy) by measuring it against the standard provided by the tacit knowledge that it attempts to specify and describe. Thus a proposed operational test for, say, segmentation into words, must meet the empirical condition of conforming, in a mass of crucial and clear cases, to the linguistic intuition of the native speaker concerning such elements. Otherwise, it is without value. The same, obviously, is true in the case of any proposed operational procedure or any proposed grammatical description. If operational procedures were available that met this test, we might be justified in relying on their results in unclear and difficult cases. This remains a hope for the future rather than a present reality, however. This is the objective situation of present-day linguistic work; allusions to presumably well-known "procedures of elicitation" or "objective methods" simply obscure the actual situation in which linguistic work must, for the present, proceed. Furthermore, there is no reason to expect that reliable operational criteria for the deeper and more important theoretical notions of linguistics (such as "grammaticalness" and "paraphrase") will ever be forthcoming.

Even though few reliable operational procedures have been developed, the theoretical (that is, grammatical) investigation of the knowledge of the native speaker can proceed perfectly well. The critical problem for grammatical theory today is not a paucity of evidence but rather the inadequacy of present theories of language to account for masses of evidence that are hardly

open to serious question. The problem for the grammarian is to construct a description and, where possible, an explanation for the enormous mass of unquestionable data concerning the linguistic intuition of the native speaker (often, himself); the problem for one concerned with operational procedures is to develop tests that give the correct results and make relevant distinctions. Neither the study of grammar nor the attempt to develop useful tests is hampered by lack of evidence with which to check results, for the present. We may hope that these efforts will converge, but they must obviously converge on the tacit knowledge of the native speaker if they are to be of any significance.

One may ask whether the necessity for present-day linguistics to give such priority to introspective evidence and to the linguistic intuition of the native speaker excludes it from the domain of science. The answer to this essentially terminological question seems to have no bearing at all on any serious issue. At most, it determines how we shall denote the kind of research that can be effectively carried out in the present state of our technique and understanding. However, this terminological question actually does relate to a different issue of some interest, namely the question whether the important feature of the successful sciences has been their search for insight or their concern for objectivity. The social and behavioral sciences provide ample evidence that objectivity can be pursued with little consequent gain in insight and understanding. On the other hand, a good case can be made for the view that the natural sciences have, by and large, sought objectivity primarily insofar as it is a tool for gaining insight (for providing phenomena that can suggest or test deeper explanatory hypotheses).

In any event, at a given stage of investigation, one whose concern is for insight and understanding (rather than for objectivity as a goal in itself) must ask whether or to what extent a wider range and more exact description of phenomena is relevant to solving the problems that he faces. In linguistics, it seems to me that sharpening of the data by more objective tests is a matter of small importance for the problems at hand. One who disagrees with this estimate of the present situation in linguistics can

justify his belief in the current importance of more objective operational tests by showing how they can lead to new and deeper understanding of linguistic structure. Perhaps the day will come when the kinds of data that we now can obtain in abundance will be insufficient to resolve deeper questions concerning the structure of language. However, many questions that can realistically and significantly be formulated today do not demand evidence of a kind that is unavailable or unattainable without significant improvements in objectivity of experimental technique.

Although there is no way to avoid the traditional assumption that the speaker-hearer's linguistic intuition is the ultimate standard that determines the accuracy of any proposed grammar, linguistic theory, or operational test, it must be emphasized, once again, that this tacit knowledge may very well not be immediately available to the user of the language. To eliminate what has seemed to some an air of paradox in this remark, let me illustrate with a few examples.

If a sentence such as "flying planes can be dangerous" is presented in an appropriately constructed context, the listener will interpret it immediately in a unique way, and will fail to detect the ambiguity. In fact, he may reject the second interpretation, when this is pointed out to him, as forced or unnatural (independently of which interpretation he originally selected under contextual pressure). Nevertheless, his intuitive knowledge of the language is clearly such that both of the interpretations (corresponding to "flying planes are dangerous" and "flying planes is dangerous") are assigned to the sentence by the grammar he has internalized in some form.

In the case just mentioned, the ambiguity may be fairly transparent. But consider such a sentence as

(5) I had a book stolen

Few hearers may be aware of the fact that their internalized grammar in fact provides at least three structural descriptions for this sentence. Nevertheless, this fact can be brought to consciousness by consideration of slight elaborations of sentence

(5), for example: (i) "I had a book stolen from my car when I stupidly left the window open," that is, "someone stole a book from my car"; (ii) "I had a book stolen from his library by a professional thief who I hired to do the job," that is, "I had someone steal a book"; (iii) "I almost had a book stolen, but they caught me leaving the library with it," that is, "I had almost succeeded in stealing a book." In bringing to consciousness the triple ambiguity of (5) in this way, we present no new information to the hearer and teach him nothing new about his language but simply arrange matters in such a way that his linguistic intuition, previously obscured, becomes evident to him.

As a final illustration, consider the sentences

(6) I persuaded John to leave

(7) I expected John to leave

The first impression of the hearer may be that these sentences receive the same structural analysis. Even fairly careful thought may fail to show him that his internalized grammar assigns very different syntactic descriptions to these sentences. In fact, so far as I have been able to discover, no English grammar has pointed out the fundamental distinction between these two constructions (in particular, my own sketches of English grammar in Chomsky, 1955, 1962a, failed to note this). However, it is clear that the sentences (6) and (7) are not parallel in structure. The difference can be brought out by consideration of the sentences

(8) (i) I persuaded a specialist to examine John

(ii) I persuaded John to be examined by a specialist

(9) (i) I expected a specialist to examine John

(ii) I expected John to be examined by a specialist

The sentences (9i) and (9ii) are "cognitively synonymous": one is true if and only if the other is true. But no variety of even weak paraphrase holds between (8i) and (8ii). Thus (8i) can be true or false quite independently of the truth or falsity of (8ii). Whatever difference of connotation or "topic" or emphasis one may find between (9i) and (9ii) is just the difference that exists be-

tween the active sentence "a specialist will examine John" and its passive counterpart "John will be examined by a specialist." This is not at all the case with respect to (8), however. In fact, the underlying deep structure for (6) and (8ii) must show that "John" is the Direct-Object of the Verb Phrase as well as the grammatical Subject of the embedded sentence. Furthermore, in (8ii) "John" is the logical Direct-Object of the embedded sentence, whereas in (8i) the phrase "a specialist" is the Direct-Object of the Verb Phrase and the logical Subject of the embedded sentence. In (7), (9i), and (9ii), however, the Noun Phrases "John," "a specialist," and "John," respectively, have no grammatical functions other than those that are internal to the embedded sentence; in particular, "John" is the logical Direct-Object and "a specialist" the logical Subject in the embedded sentences of (9). Thus the underlying deep structures for (8i), (8ii), (9i), and (9ii) are, respectively, the following:¹³

- (10) (i) Noun Phrase — Verb — Noun Phrase — Sentence
(I — persuaded — a specialist — a specialist will examine John)
- (ii) Noun Phrase — Verb — Noun Phrase — Sentence
(I — persuaded — John — a specialist will examine John)
- (11) (i) Noun Phrase — Verb — Sentence
(I — expected — a specialist will examine John)
- (ii) Noun Phrase — Verb — Sentence
(I — expected — a specialist will examine John)

In the case of (10ii) and (11ii), the passive transformation will apply to the embedded sentence, and in all four cases other operations will give the final surface forms of (8) and (9). The important point in the present connection is that (8i) differs from (8ii) in underlying structure, although (9i) and (9ii) are essentially the same in underlying structure. This accounts for the difference in meaning. Notice, in support of this difference in analysis, that we can have "I persuaded John that (of the fact that) Sentence," but not "I expected John that (of the fact that) Sentence."

The example (6)–(7) serves to illustrate two important points. First, it shows how unrevealing surface structure may be as to underlying deep structure. Thus (6) and (7) are the same in surface structure, but very different in the deep structure that underlies them and determines their semantic interpretations. Second, it illustrates the elusiveness of the speaker's tacit knowledge. Until such examples as (8) and (9) are adduced, it may not be in the least clear to a speaker of English that the grammar that he has internalized in fact assigns very different syntactic analyses to the superficially analogous sentences (6) and (7).

In short, we must be careful not to overlook the fact that surface similarities may hide underlying distinctions of a fundamental nature, and that it may be necessary to guide and draw out the speaker's intuition in perhaps fairly subtle ways before we can determine what is the actual character of his knowledge of his language or of anything else. Neither point is new (the former is a commonplace of traditional linguistic theory and analytic philosophy; the latter is as old as Plato's *Meno*); both are too often overlooked.

A grammar can be regarded as a theory of a language; it is *descriptively adequate* to the extent that it correctly describes the intrinsic competence of the idealized native speaker. The structural descriptions assigned to sentences by the grammar, the distinctions that it makes between well-formed and deviant, and so on, must, for descriptive adequacy, correspond to the linguistic intuition of the native speaker (whether or not he may be immediately aware of this) in a substantial and significant class of crucial cases.

A linguistic theory must contain a definition of "grammar," that is, a specification of the class of potential grammars. We may, correspondingly, say that a *linguistic theory is descriptively adequate* if it makes a descriptively adequate grammar available for each natural language.

Although even descriptive adequacy on a large scale is by no means easy to approach, it is crucial for the productive development of linguistic theory that much higher goals than this be pursued. To facilitate the clear formulation of deeper questions,

it is useful to consider the abstract problem of constructing an "acquisition model" for language, that is, a theory of language learning or grammar construction. Clearly, a child who has learned a language has developed an internal representation of a system of rules that determine how sentences are to be formed, used, and understood. Using the term "grammar" with a systematic ambiguity (to refer, first, to the native speaker's internally represented "theory of his language" and, second, to the linguist's account of this), we can say that the child has developed and internally represented a generative grammar, in the sense described. He has done this on the basis of observation of what we may call *primary linguistic data*. This must include examples of linguistic performance that are taken to be well-formed sentences, and may include also examples designated as non-sentences, and no doubt much other information of the sort that is required for language learning, whatever this may be (see pp. 31–32). On the basis of such data, the child constructs a grammar — that is, a theory of the language of which the well-formed sentences of the primary linguistic data constitute a small sample.¹⁴ To learn a language, then, the child must have a method for devising an appropriate grammar, given primary linguistic data. As a precondition for language learning, he must possess, first, a linguistic theory that specifies the form of the grammar of a possible human language, and, second, a strategy for selecting a grammar of the appropriate form that is compatible with the primary linguistic data. As a long-range task for general linguistics, we might set the problem of developing an account of this innate linguistic theory that provides the basis for language learning. (Note that we are again using the term "theory" — in this case "theory of language" rather than "theory of a particular language" — with a systematic ambiguity, to refer both to the child's innate predisposition to learn a language of a certain type and to the linguist's account of this.)

To the extent that a linguistic theory succeeds in selecting a descriptively adequate grammar on the basis of primary linguistic data, we can say that it meets the condition of *explanatory adequacy*. That is, to this extent, it offers an explanation for the

intuition of the native speaker on the basis of an empirical hypothesis concerning the innate predisposition of the child to develop a certain kind of theory to deal with the evidence presented to him. Any such hypothesis can be falsified (all too easily, in actual fact) by showing that it fails to provide a descriptively adequate grammar for primary linguistic data from some other language—evidently the child is not predisposed to learn one language rather than another. It is supported when it does provide an adequate explanation for some aspect of linguistic structure, an account of the way in which such knowledge might have been obtained.

Clearly, it would be utopian to expect to achieve explanatory adequacy on a large scale in the present state of linguistics. Nevertheless, considerations of explanatory adequacy are often critical for advancing linguistic theory. Gross coverage of a large mass of data can often be attained by conflicting theories; for precisely this reason it is not, in itself, an achievement of any particular theoretical interest or importance. As in any other field, the important problem in linguistics is to discover a complex of data that differentiates between conflicting conceptions of linguistic structure in that one of these conflicting theories can describe these data only by *ad hoc* means whereas the other can explain it on the basis of some empirical assumption about the form of language. Such small-scale studies of explanatory adequacy have, in fact, provided most of the evidence that has any serious bearing on the nature of linguistic structure. Thus whether we are comparing radically different theories of grammar or trying to determine the correctness of some particular aspect of one such theory, it is questions of explanatory adequacy that must, quite often, bear the burden of justification. This remark is in no way inconsistent with the fact that explanatory adequacy on a large scale is out of reach, for the present. It simply brings out the highly tentative character of any attempt to justify an empirical claim about linguistic structure.

To summarize briefly, there are two respects in which one can speak of "justifying a generative grammar." On one level (that

of descriptive adequacy), the grammar is justified to the extent that it correctly describes its object, namely the linguistic intuition—the tacit competence—of the native speaker. In this sense, the grammar is justified on *external* grounds, on grounds of correspondence to linguistic fact. On a much deeper and hence much more rarely attainable level (that of explanatory adequacy), a grammar is justified to the extent that it is a *principled* descriptively adequate system, in that the linguistic theory with which it is associated selects this grammar over others, given primary linguistic data with which all are compatible. In this sense, the grammar is justified on *internal* grounds, on grounds of its relation to a linguistic theory that constitutes an explanatory hypothesis about the form of language as such. The problem of internal justification—of explanatory adequacy—is essentially the problem of constructing a theory of language acquisition, an account of the specific innate abilities that make this achievement possible.

§ 5. FORMAL AND SUBSTANTIVE UNIVERSALS

A theory of linguistic structure that aims for explanatory adequacy incorporates an account of linguistic universals, and it attributes tacit knowledge of these universals to the child. It proposes, then, that the child approaches the data with the presumption that they are drawn from a language of a certain antecedently well-defined type, his problem being to determine which of the (humanly) possible languages is that of the community in which he is placed. Language learning would be impossible unless this were the case. The important question is: What are the initial assumptions concerning the nature of language that the child brings to language learning, and how detailed and specific is the innate schema (the general definition of "grammar") that gradually becomes more explicit and differentiated as the child learns the language? For the present we cannot come at all close to making a hypothesis about innate schemata that is rich, detailed, and specific enough to account for the fact of language acquisition. Consequently, the main

task of linguistic theory must be to develop an account of linguistic universals that, on the one hand, will not be falsified by the actual diversity of languages and, on the other, will be sufficiently rich and explicit to account for the rapidity and uniformity of language learning, and the remarkable complexity and range of the generative grammars that are the product of language learning.

The study of linguistic universals is the study of the properties of any generative grammar for a natural language. Particular assumptions about linguistic universals may pertain to either the syntactic, semantic, or phonological component, or to interrelations among the three components.

It is useful to classify linguistic universals as *formal* or *substantive*. A theory of substantive universals claims that items of a particular kind in any language must be drawn from a fixed class of items. For example, Jakobson's theory of distinctive features can be interpreted as making an assertion about substantive universals with respect to the phonological component of a generative grammar. It asserts that each output of this component consists of elements that are characterized in terms of some small number of fixed, universal, phonetic features (perhaps on the order of fifteen or twenty), each of which has a substantive acoustic-articulatory characterization independent of any particular language. Traditional universal grammar was also a theory of substantive universals, in this sense. It not only put forth interesting views as to the nature of universal phonetics, but also advanced the position that certain fixed syntactic categories (Noun, Verb, etc.) can be found in the syntactic representations of the sentences of any language, and that these provide the general underlying syntactic structure of each language. A theory of substantive semantic universals might hold for example, that certain designative functions must be carried out in a specified way in each language. Thus it might assert that each language will contain terms that designate persons or lexical items referring to certain specific kinds of objects, feelings, behavior, and so on.

It is also possible, however, to search for universal properties

of a more abstract sort. Consider a claim that the grammar of every language meets certain specified formal conditions. The truth of this hypothesis would not in itself imply that any particular rule must appear in all or even in any two grammars. The property of having a grammar meeting a certain abstract condition might be called a *formal* linguistic universal, if shown to be a general property of natural languages. Recent attempts to specify the abstract conditions that a generative grammar must meet have produced a variety of proposals concerning formal universals, in this sense. For example, consider the proposal that the syntactic component of a grammar must contain transformational rules (these being operations of a highly special kind) mapping semantically interpreted deep structures into phonetically interpreted surface structures, or the proposal that the phonological component of a grammar consists of a sequence of rules, a subset of which may apply cyclically to successively more dominant constituents of the surface structure (a transformational cycle, in the sense of much recent work on phonology). Such proposals make claims of a quite different sort from the claim that certain substantive phonetic elements are available for phonetic representation in all languages, or that certain specific categories must be central to the syntax of all languages, or that certain semantic features or categories provide a universal framework for semantic description. Substantive universals such as these concern the vocabulary for the description of language; formal universals involve rather the character of the rules that appear in grammars and the ways in which they can be interconnected.

On the semantic level, too, it is possible to search for what might be called formal universals, in essentially the sense just described. Consider, for example, the assumption that proper names, in any language, must designate objects meeting a condition of spatiotemporal contiguity,¹⁵ and that the same is true of other terms designating objects; or the condition that the color words of any language must subdivide the color spectrum into continuous segments; or the condition that artifacts are defined in terms of certain human goals, needs, and functions instead of solely in terms of physical qualities.¹⁶ Formal con-

straints of this sort on a system of concepts may severely limit the choice (by the child, or the linguist) of a descriptive grammar, given primary linguistic data.

The existence of deep-seated formal universals, in the sense suggested by such examples as these, implies that all languages are cut to the same pattern, but does not imply that there is any point by point correspondence between particular languages. It does not, for example, imply that there must be some reasonable procedure for translating between languages.¹⁷

In general, there is no doubt that a theory of language, regarded as a hypothesis about the innate "language-forming capacity" of humans, should concern itself with both substantive and formal universals. But whereas substantive universals have been the traditional concern of general linguistic theory, investigations of the abstract conditions that must be satisfied by any generative grammar have been undertaken only quite recently. They seem to offer extremely rich and varied possibilities for study in all aspects of grammar.

§ 6. FURTHER REMARKS ON DESCRIPTIVE AND EXPLANATORY THEORIES

Let us consider with somewhat greater care just what is involved in the construction of an "acquisition model" for language. A child who is capable of language learning must have

- (12) (i) a technique for representing input signals
- (ii) a way of representing structural information about these signals
- (iii) some initial delimitation of a class of possible hypotheses about language structure
- (iv) a method for determining what each such hypothesis implies with respect to each sentence
- (v) a method for selecting one of the (presumably, infinitely many) hypotheses that are allowed by (iii) and are compatible with the given primary linguistic data

Correspondingly, a theory of linguistic structure that aims for explanatory adequacy must contain

- (13) (i) a universal phonetic theory that defines the notion "possible sentence"
- (ii) a definition of "structural description"
- (iii) a definition of "generative grammar"
- (iv) a method for determining the structural description of a sentence, given a grammar
- (v) a way of evaluating alternative proposed grammars

Putting the same requirements in somewhat different terms, we must require of such a linguistic theory that it provide for

- (14) (i) an enumeration of the class s_1, s_2, \dots of possible sentences
- (ii) an enumeration of the class SD_1, SD_2, \dots of possible structural descriptions
- (iii) an enumeration of the class G_1, G_2, \dots of possible generative grammars
- (iv) specification of a function f such that $SD_{f(i,j)}$ is the structural description assigned to sentence s_i by grammar G_j , for arbitrary i, j ¹⁸
- (v) specification of a function m such that $m(i)$ is an integer associated with the grammar G_i as its value (with, let us say, lower value indicated by higher number)

Conditions of at least this strength are entailed by the decision to aim for explanatory adequacy.

A theory meeting these conditions would attempt to account for language learning in the following way. Consider first the nature of primary linguistic data. This consists of a finite amount of information about sentences, which, furthermore, must be rather restricted in scope, considering the time limitations that are in effect, and fairly degenerate in quality (cf. note 14). For example, certain signals might be accepted as properly formed sentences, while others are classed as nonsentences, as a result of correction of the learner's attempts on the part of the linguistic community. Furthermore, the conditions of use might be such