

Chomsky, N. (1986). *Knowledge of language: Its nature, origin and use*. New York: Praeger.

1

Knowledge of Language as a Focus of Inquiry

The study of language has a long and rich history, extending over thousands of years. This study has frequently been understood as an inquiry into the nature of mind and thought on the assumption that "languages are the best mirror of the human mind" (Leibniz). A common conception was that "with respect to its *substance* grammar is one and the same in all languages, though it does vary *accidentally*" (Roger Bacon). The invariant "substance" was often taken to be the mind and its acts; particular languages use various mechanisms—some rooted in human reason, others arbitrary and adventitious—for the expression of thought, which is a constant across languages. One leading eighteenth century rational grammarian defined "general grammar" as a deductive science concerned with "the immutable and general principles of spoken or written language" and their consequences; it is "prior to all languages," because its principles "are the same as those that direct human reason in its intellectual operations" (Beauzée). Thus, "the science of language does not differ at all from the science of thought." "Particular grammar" is not a true "science" in the sense of this rationalist tradition because it is not based solely on universal necessary laws; it is an "art" or technique that shows how given languages realize the general principles of human reason. As John Stuart Mill later expressed the same leading idea, "The principles and rules of grammar are the means by which the forms of language are made to correspond with the universal forms of thought. . . . The structure of every sentence is a lesson in logic." Others, particularly during the

Romantic period, argued that the nature and content of thought are determined in part by the devices made available for its expression in particular languages. These devices may include contributions of individual genius that affect the "character" of a language, enriching its means of expression and the thoughts expressed without affecting its "form," its sound system and rules of word and sentence formation (Humboldt).

With regard to the acquisition of knowledge, it was widely held that the mind is not "so much to be filled therewith from without, like a vessel, as to be kindled and awaked" (Ralph Cudworth); "The growth of knowledge...[rather resembles]... the growth of Fruit; however external causes may in some degree cooperate, it is the internal vigour, and virtue of the tree, that must ripen the juices to their just maturity" (James Harris).¹ Applied to language, this essentially Platonistic conception would suggest that knowledge of a particular language grows and matures along a course that is in part intrinsically determined, with modifications reflecting observed usage, rather in the manner of the visual system or other bodily "organs" that develop along a course determined by genetic instructions under the triggering and shaping effects of environmental factors.

With the exception of the relativism of the Romantics, such ideas were generally regarded with much disapproval in the mainstream of linguistic research by the late nineteenth century and on through the 1950s. In part, this attitude developed under the impact of a rather narrowly construed empiricism and later behaviorist and operationalist doctrine. In part, it resulted from the quite real and impressive successes of historical and descriptive studies conducted within a narrower compass, specifically, the discovery of "sound laws" that provided much understanding of the history of languages and their relationships. In part, it was a natural consequence of the investigation of a much richer variety of languages than were known to earlier scholars, languages that appeared to violate many of the allegedly *a priori* conceptions of the earlier rationalist tradition.² After a century of general neglect or obloquy, ideas resembling those of the earlier tradition re-emerged (initially, with virtually no awareness of historical antecedents) in the mid-1950s, with the development of what came to be called "generative grammar"—again, reviving a long-lapsed and largely forgotten tradition.³

The generative grammar of a particular language (where "generative" means nothing more than "explicit") is a theory that is concerned with the form and meaning of expressions of this language. One can imagine many different kinds of approach to such questions, many points of view that might be adopted in dealing with them. Generative grammar limits itself to certain elements of this larger picture. Its standpoint is that of individual psychology. It is concerned with those aspects of form and meaning that are determined by the "language faculty," which is understood to be a particular component of the human mind. The nature of this faculty is the subject matter of a general theory of linguistic structure that aims to discover the framework of principles and elements common to attainable human languages; this theory is now often called "universal grammar" (UG), adapting a traditional term to a new context of inquiry. UG may be regarded as a characterization of the genetically determined language faculty. One may think of this faculty as a "language acquisition device," an innate component of the human mind that yields a particular language through interaction with presented experience, a device that converts experience into a system of knowledge attained: knowledge of one or another language.

The study of generative grammar represented a significant shift of focus in the approach to problems of language. Put in the simplest terms, to be elaborated below, the shift of focus was from behavior or the products of behavior to states of the mind/brain that enter into behavior. If one chooses to focus attention on this latter topic, the central concern becomes knowledge of language: its nature, origins, and use.

The three basic questions that arise, then, are these:

- (i) What constitutes knowledge of language? (1)
- (ii) How is knowledge of language acquired?
- (iii) How is knowledge of language put to use?

The answer to the first question is given by a particular generative grammar, a theory concerned with the state of the mind/brain of the person who knows a particular language. The answer to the second is given by a specification of UG along with an account of the ways in which its principles interact with experience to yield a particular language; UG is a theory of the "initial state" of the language faculty, prior to any

linguistic experience. The answer to the third question would be a theory of how the knowledge of language attained enters into the expression of thought and the understanding of presented specimens of language, and derivatively, into communication and other special uses of language.

So far, this is nothing more than the outline of a research program that takes up classical questions that had been put aside for many years. As just described, it should not be particularly controversial, since it merely expresses an interest in certain problems and offers a preliminary analysis of how they might be confronted, although as is often the case, the initial formulation of a problem may prove to be far-reaching in its implications, and ultimately controversial as it is developed.

Some elements of this picture may appear to be more controversial than they really are. Consider, for example, the idea that there is a language faculty, a component of the mind/brain that yields knowledge of language given presented experience. It is not at issue that humans attain knowledge of English, Japanese, and so forth, while rocks, birds, or apes do not under the same (or indeed any) conditions. There is, then, some property of the mind/brain that differentiates humans from rocks, birds, or apes. Is this a distinct "language faculty" with specific structure and properties, or, as some believe, is it the case that humans acquire language merely by applying generalized learning mechanisms of some sort, perhaps with greater efficiency or scope than other organisms? These are not topics for speculation or *a priori* reasoning but for empirical inquiry, and it is clear enough how to proceed: namely, by facing the questions of (1). We try to determine what is the system of knowledge that has been attained and what properties must be attributed to the initial state of the mind/brain to account for its attainment. Insofar as these properties are language-specific, either individually or in the way they are organized and composed, there is a distinct language faculty.

Generative grammar is sometimes referred to as a theory, advocated by this or that person. In fact, it is not a theory any more than chemistry is a theory. Generative grammar is a topic, which one may or may not choose to study. Of course, one can adopt a point of view from which chemistry disappears as a discipline (perhaps it is all done by angels with mirrors). In this sense, a decision to study chemistry does stake out a position on

matters of fact. Similarly, one may argue that the topic of generative grammar does not exist, although it is hard to see how to make this position minimally plausible. Within the study of generative grammar there have been many changes and differences of opinion, often reversion to ideas that had been abandoned and were later reconstructed in a different light. Evidently, this is a healthy phenomenon indicating that the discipline is alive, although it is sometimes, oddly, regarded as a serious deficiency, a sign that something is wrong with the basic approach. I will review some of these changes as we proceed.

In the mid-1950s, certain proposals were advanced as to the form that answers to the questions of (1) might take, and a research program was inaugurated to investigate the adequacy of these proposals and to sharpen and apply them. This program was one of the strands that led to the development of the cognitive sciences in the contemporary sense, sharing with other approaches the belief that certain aspects of the mind/brain can be usefully construed on the model of computational systems of rules that form and modify representations, and that are put to use in interpretation and action. From its origins (or with a longer perspective, one might say "its reincarnation") about 30 years ago, the study of generative grammar was undertaken with an eye to gaining some insight into the nature and origins of systems of knowledge, belief, and understanding more broadly, in the hope that these general questions could be illuminated by a detailed investigation of the special case of human language.

This research program has since been running its course, along a number of different paths. I will be concerned here with only one of these, with the problems it faced and the steps that were taken in an effort to deal with them. During the past 5–6 years, these efforts have converged in a somewhat unexpected way, yielding a rather different conception of the nature of language and its mental representation, one that offers interesting answers to a range of empirical questions and opens a variety of new ones to inquiry while suggesting a rethinking of the character of others. This is what accounts for an unmistakable sense of energy and anticipation—and also uncertainty—which is reminiscent of the period when the study of generative grammar in the modern sense was initiated about 30 years ago.

Some of the work now being done is quite different in character from what had previously been possible as well as considerably broader in empirical scope, and it may be that results of a rather new kind are within reach, or at least within sight. I would like to try to explain why this may be so, beginning with some remarks about goals, achievements, and failures of the past years.

To avoid misunderstanding, I am not speaking here about all of the study of language but rather of generative grammar, and even here I will not attempt anything like a real history of the course of research but rather will give a somewhat idealized picture that is in part clearer in retrospect than it was at the time. Furthermore, what I am describing has represented a minority position throughout, and probably still does, although in my view it is the correct one. A number of different current approaches share properties of the sort discussed here and may be intertranslatable to a considerable extent. I will not consider this important topic here and will also make no effort to survey the range of ideas, often conflicting, that fall within the particular tendency that I will discuss—what is now sometimes called “government-binding (GB) theory.”

I want to consider, then, two major conceptual shifts, one that inaugurated the contemporary study of generative grammar, and a second, more theory-internal, that is now in process and that offers some new perspectives on traditional problems.⁴

Traditional and structuralist grammar did not deal with the questions of (1), the former because of its implicit reliance on the unanalyzed intelligence of the reader, the latter because of its narrowness of scope. The concerns of traditional and generative grammar are, in a certain sense, complementary: a good traditional or pedagogical grammar provides a full list of exceptions (irregular verbs, etc.), paradigms and examples of regular constructions, and observations at various levels of detail and generality about the form and meaning of expressions. But it does not examine the question of how the reader of the grammar uses such information to attain the knowledge that is used to form and interpret new expressions, or the question of the nature and elements of this knowledge: essentially the questions of (1), above. Without too much exaggeration, one could describe such a grammar as a structured and organized version of the data presented to a child learning a language,

with some general commentary and often insightful observations. Generative grammar, in contrast, is concerned primarily with the intelligence of the reader, the principles and procedures brought to bear to attain full knowledge of a language. Structuralist theories, both in the European and American traditions, did concern themselves with analytic procedures for deriving aspects of grammar from data, as in the procedural theories of Nikolay Trubetzkoy, Zellig Harris, Bernard Bloch, and others, but primarily in the areas of phonology and morphology. The procedures suggested were seriously inadequate and in any event could not possibly be understood (and were not intended) to provide an answer to question (1ii), even in the narrower domains where most work was concentrated. Nor was there an effort to determine what was involved in offering a comprehensive account of the knowledge of the speaker/hearer.

As soon as these questions were squarely faced, a wide range of new phenomena were discovered, including quite simple ones that had passed unnoticed, and severe problems arose that had previously been ignored or seriously misunderstood. A standard belief 30 years ago was that language acquisition is a case of "overlearning." Language was regarded as a habit system, one that was assumed to be much overdetermined by available evidence. Production and interpretation of new forms was taken to be a straightforward matter of analogy, posing no problems of principle.⁵ Attention to the questions of (1) quickly reveals that exactly the opposite is the case: language poses in a sharp and clear form what has sometimes been called "Plato's problem," the problem of "poverty of stimulus," of accounting for the richness, complexity, and specificity of shared knowledge, given the limitations of the data available. This difference of perception concerning where the problem lies—overlearning or poverty of evidence—reflects very clearly the effect of the shift of focus that inaugurated the study of generative grammar.

A great many examples have been given over the years to illustrate what clearly is the fundamental problem: the problem of poverty of evidence. A familiar example is the structure-dependence of rules, the fact that without instruction or direct evidence, children unerringly use computationally complex structure-dependent rules rather than computationally simple rules that involve only the predicate "leftmost" in a linear

sequence of words.⁶ To take some other examples, to which we will return, consider sentences (2)–(7):

- | | |
|---|-----|
| I wonder who [the men expected to see them] | (2) |
| [the men expected to see them] | (3) |
| John ate an apple | (4) |
| John ate | (5) |
| John is too stubborn to talk to Bill | (6) |
| John is too stubborn to talk to | (7) |

Both (2) and (3) include the clause bounded by brackets, but only in (2) may the pronoun *them* be referentially dependent on the antecedent *the men*; in (3) the pronoun is understood as referring in some manner indicated in the situational or discourse context, but not to the men. Numerous facts of this sort, falling under what is now generally called “binding theory,” are known without relevant experience to differentiate the cases. Such facts pose a serious problem that was not recognized in earlier work: How does every child know, unerringly, to interpret the clause differently in the two cases? And why does no pedagogic grammar have to draw the learner’s attention to such facts (which were, in fact, noticed only quite recently, in the course of the study of explicit rule systems in generative grammar)?

Turning to examples (4)–(7), sentence (5) means that John ate something or other, a fact that one might explain on the basis of a simple inductive procedure: *ate* takes an object, as in (4), and if the object is missing, it is understood as arbitrary. Applying the same inductive procedure to (6) and (7), it should be that (7) means that John is so stubborn that he (John) will not talk to some arbitrary person, on the analogy of (6). But the meaning is, in fact, quite different: namely, that John is so stubborn that some arbitrary person won’t talk to him (John). Again, this is known without training or relevant evidence.⁷

The situation is, in fact, more complex. Although plausible, the inductive procedure suggested for the relatively straightforward examples (4)–(5) does not seem correct. As noted by Howard Lasnik, the word *eat* has a somewhat different meaning in its intransitive usage, something like *dine*. One can say “John ate his shoe,” but “John ate” cannot be understood to include this case. The observation is general for such cases. The

intransitive forms differ from normal intransitives in other respects; for example, we can form "the dancing bear" (corresponding to "the bear that dances"), but not "the eating man" (corresponding to "the man who eats").⁸ Such facts pose further problems of poverty of stimulus.

Children do not make errors about the interpretation of such sentences as (6)–(7) past a certain stage of development, and if they did, the errors would largely be uncorrectable. It is doubtful that even the most compendious traditional or teaching grammar notes such simple facts as those illustrated in (2)–(7), and such observations lie far beyond the domain of structural grammars. A wide variety of examples of this sort immediately come to attention when one faces the questions formulated in (1).

Knowledge of language is often characterized as a practical ability to speak and understand, so that questions (li) and (liii) are closely related, perhaps identified. Ordinary usage makes a much sharper distinction between the two questions, and is right to do so. Two people may share exactly the same knowledge of language but differ markedly in their ability to put this knowledge to use. Ability to use language may improve or decline without any change in knowledge. This ability may also be impaired, selectively or in general, with no loss of knowledge, a fact that would become clear if injury leading to impairment recedes and lost ability is recovered. Many such considerations support the commonsense assumption that knowledge cannot be properly described as a practical ability. Furthermore, even if this view could somehow be maintained, it would leave open all of the serious questions. Thus, what is the nature of the "practical ability" manifested in our interpretation of the sentences (2)–(7), how is it properly described, and how is it acquired?

Often it is not immediately obvious what our knowledge of language entails in particular cases, a fact illustrated even with short and simple sentences such as (8)–(10):

- his wife loves her husband (8)
- John is too clever to expect us to catch Bill (9)
- John is too clever to expect us to catch (10)

In the case of (8), it takes some thought to determine whether *his* can be referentially dependent on *her husband* if *her* is

dependent on *his wife*—that is, if the reference of either *he* or *she* is not somehow contextually indicated.⁹ Examples (9) and (10) are, in fact, analogous to (6) and (7), respectively, but again, it takes some thought to discover that (10) means that John is so clever that an arbitrary person cannot expect us to catch him (John), although it is clear at once that it does not mean that John is so clever that he (John) cannot catch some arbitrary person, on the analogy of (9) (and (4), (5)). Our abilities seem limited somehow in such cases (and there are far more complex ones), but it would make little sense to speak of our knowledge of language as “limited” in any comparable way.

Suppose we insist on speaking of knowledge of language, as a practical ability to speak and understand. Then normal usage must be revised in numerous cases such as those just discussed. Suppose that Jones takes a public speaking course and improves his ability to speak and understand without any change in his knowledge of English, as we would describe the situation in normal usage. We must now revise this common-sense usage and say, rather, that Jones has improved his ability₁ to use his ability₂ to speak and understand; similar translations are required in the other cases. But the two occurrences of “ability” in this description are hardly more than homonyms. Ability₁ is ability in the normal sense of the word: it can improve or decline, can be inadequate to determine consequences of knowledge, and so on. Ability₂, however, remains stable while our ability to use it changes, and we have this kind of “ability” even when we are unable to detect what it entails in concrete cases. In short, the neologism “ability₂” is invested with all the properties of knowledge. Note that there are cases when we do speak of abilities that we cannot put to use: for example, the case of swimmers who cannot swim because their hands are tied, although they retain the ability to swim. The cases in question are not of this sort, however.

The purpose of the attempt to reduce knowledge to ability is, presumably, to avoid problematic features that seem to inhere in the concept of knowledge, to show that these can be explained in dispositional or other terms more closely related to actual behavior (whether this is possible even in the case of ability₁, the normal sense, is another question). But nothing of the sort is achieved by this departure from ordinary usage; the

problems remain, exactly as before, now embedded in terminological confusion. The task of determining the nature of our knowledge (= ability₂), and accounting for its origins and use, remains exactly as challenging as before, despite the terminological innovations.

Other examples similar to (8)–(10) raise further questions. Consider the following sentences:

John is too stubborn to expect anyone to talk to Bill (11)

John is too stubborn to visit anyone who talked to Bill (12)

Suppose we delete *Bill* from (11) and (12), yielding (13) and (14), respectively:

John is too stubborn to expect anyone to talk to (13)

John is too stubborn to visit anyone who talked to (14)

Sentence (13) is structurally analogous to (10), and is understood in the same manner: it means that John is so stubborn that an arbitrary person would not expect anyone to talk to him (John). "By analogy," then, we would expect sentence (14) to mean that John is so stubborn that an arbitrary person would not visit anyone who talked to him (John). But it does not have that meaning; in fact, it is gibberish. Here we have a double failure of analogy. Sentence (14) is not understood "on the analogy" of (4), (5), (6), (9), and (12) (hence meaning that John is so stubborn that he (John) would not visit anyone who talked to some arbitrary person), nor is it understood "on the analogy" of (7), (10), and (13); rather, it has no interpretation at all. And while the status of (11), (12), and (14) is immediately obvious, it takes some thought or preparation to see that (13) has the interpretation it does have, and thus to determine the consequences of our knowledge in this case.

Again, these are facts that we know, however difficult it may be to determine that our system of knowledge has these consequences. We know these facts without instruction or even

direct evidence, surely without correction of error by the speech community. It would be absurd to try to teach such facts as these to people learning English as a second language, just as no one taught them to us or even presented us with evidence that could yield this knowledge by any generally reliable procedure. This is knowledge without grounds, without good reasons or support by reliable procedures in any general or otherwise useful sense of these notions. Were we to insist that knowledge is a kind of ability, we would have to claim that we lack the ability to understand "John is too stubborn to talk to" as meaning "John is too stubborn to talk to someone or other" (on the analogy of "John ate an apple" — "John ate"), and that we lack the ability to understand (14) on the analogy of "John ate an apple" — "John ate" (so that it means that John is too stubborn to visit anyone who talked to someone or other) or on the analogy of "John is too stubborn to talk to," with the "inversion strategy" that we somehow use in this case (so that (14) means that John is too stubborn for someone or other to visit anyone who talked to him, John). But these would be odd claims, to say the least. These are not failures of ability. It is not that we are too weak, or lack some special skill that could be acquired. We are perfectly capable of associating the sentence (14), for example, with either of the two meanings that would be provided "by analogy" (or others), but we know that these are not the associations that our knowledge of the language provides; ability is one thing, knowledge something quite different. The system of knowledge that has somehow developed in our minds has certain consequences, not others; it relates sound and meaning and assigns structural properties to physical events in certain ways, not others.

It seems that there is little hope in accounting for our knowledge in terms of such ideas as analogy, induction, association, reliable procedures, good reasons, and justification in any generally useful sense, or in terms of "generalized learning mechanisms" (if such exist). And it seems that we should follow normal usage in distinguishing clearly between knowledge and ability to use that knowledge. We should, so it appears, think of knowledge of language as a certain state of the mind/brain, a relatively stable element in transitory mental states once it is attained; furthermore, as a state of some distinguishable

faculty of the mind—the language faculty—with its specific properties, structure, and organization, one “module” of the mind.¹⁰

NOTES

1. On these and many other discussions, primarily in the seventeenth–nineteenth centuries, see Chomsky (1966). For discussion of some misinterpretation of this work, see Bracken (1984).

2. The alleged *a priorism* of work in this tradition has often been exaggerated. See Chomsky (1966) and more recent work for discussion of this point.

3. The tradition, in this case, is a different one, represented in its most advanced form in the early work of the Indian grammarians 2,500 years ago. See Kiparsky (1982). A modern counterpart is Bloomfield (1939), which was radically different in character from the work of the period and inconsistent with his own theories of language, and remained virtually without influence or even awareness despite Bloomfield's great prestige.

4. See Newmeyer (1980) for one view of the history of this period prior to the second major conceptual shift; and for some more personal comments, the introduction to Chomsky (1975a), a somewhat abbreviated version of a 1956 revision of a 1955 manuscript, both unpublished. See Lightfoot (1982) and Hornstein and Lightfoot (1981) for discussion of the general backgrounds for much current work, and Radford (1981) for an introduction to the work that led to the second conceptual shift. See Chomsky (1981) for a more technical presentation of some of the ideas that entered into this conceptual shift and van Riemsdijk and Williams (1985) for an introductory study of this current work.

5. Although basically adopting this point of view, W.V. Quine, however, argued that there is a very severe, in fact, insuperable problem of underdetermination affecting all aspects of language and grammar, and much of psychology more generally (Quine, 1960, 1972). I do not think that he succeeded in showing that some novel form of indeterminacy affects the study of language beyond the normal underdetermination of theory by evidence; his own formulations of the thesis furthermore involve internal inconsistency (see Chomsky, 1975b, 1980b). There seems no reason on these grounds, then, to distinguish linguistics or psychology in principle from the natural sciences in accordance with what Hockney (1975) calls Quine's “bifurcation

thesis." A similar conclusion is reached by Putnam (1981) in his abandonment of metaphysical realism on Quinean grounds. His step also abandons the bifurcation thesis, although in the opposite direction.

6. See Chomsky (1975a). See Crain and Nakayama (1984) for empirical study of this question with 3–5-year-old children.

7. The reaction to such phenomena, also unnoticed until recently, again illustrates the difference of outlook of structuralist-descriptive and generative grammar. For some practitioners of the former, the statement of the facts, which is straightforward enough once they are observed, is the answer—nothing else is necessary; for the latter, the statement of the facts poses the problem to be solved. Cf. Ney (1983), particularly, his puzzlement about the "peculiar view of grammar [that] unnecessarily complicates the whole matter" by seeking an explanation for the facts. Note that there is no question of right or wrong here, but rather of topic of inquiry.

8. In early work, such facts were used to motivate an analysis of intransitives such as *eat* as derived from corresponding transitives by a system of ordered rules that excluded the unwanted cases; see Chomsky (1962).

9. On structures of this type, and problems of binding theory, more generally, see Higginbotham (1988a), among much other work.

10. See Fodor (1983). But it is too narrow to regard the "language module" as an input system in Fodor's sense, if only because it is used in speaking and thought. We might consider supplementing this picture by adding an "output system," but plainly this must be linked to the input system; we do not expect a person to speak only English and understand only Japanese. That is, the input and output systems must each access a fixed system of knowledge. The latter, however, is a central system which has essential problems of modularity, a fact that brings the entire picture into question. Furthermore, even regarded as an input system, the language module does not appear to have the property of rapidity of access that Fodor discusses, as indicated by (8)–(14). Note also that even if Fodor is right in believing that there is a sharp distinction between modules in his sense and "the rest," which is holistic in several respects, it does not follow that the residue is unstructured. In fact, this seems highly unlikely, if only because of the "epistemic boundedness" that he notes. Many other questions arise concerning Fodor's very intriguing discussion of these issues, which I will not pursue here.