23 On the Nature, Use, and Acquisition of Language

Noam Chomsky

For about thirty years, the study of language—or more accurately, one substantial component of it—has been conducted within a framework that understands linguistics to be a part of psychology, ultimately human biology. This approach attempts to reintroduce into the study of language several concerns that have been central to Western thought for thousands of years, and that have deep roots in other traditions as well: questions about the nature and origin of knowledge in particular. This approach has also been concerned to assimilate the study of language to the main body of the natural sciences. This meant, in the first place, abandoning dogmas that are entirely foreign to the natural sciences and that have no place in rational inquiry, the dogmas of the several varieties of behaviorism, for example, which seek to impose a priori limits on possible theory construction, a conception that would properly be dismissed as entirely irrational in the natural sciences. It means a frank adherence to mentalism, where we understand talk about the mind to be talk about the brain at an abstract level at which, so we try to demonstrate, principles can be formulated that enter into successful and insightful explanation of linguistic (and other) phenomena that are provided by observation and experiment. Mentalism, in this sense, has no taint of mysticism and carries no dubious ontological burden. Rather, mentalism falls strictly within the standard practice of the natural sciences and, in fact, is nothing other than the approach of the natural sciences applied to this particular domain. This conclusion, which is the opposite of what is often assumed, becomes understandable and clear if we consider specific topics in the natural sciences; for example, nineteenth-century chemistry, which sought to explain phenomena in terms of such abstract notions as elements, the periodic table, valence, benzene rings, and so on—that is, in terms of abstract

From N. Chomsky, On the nature, use and acquisition of language, in N. Chomsky, *Generative grammar: Its basis, development and prospects* (1987). Kyoto: Kyoto University of Foreign Studies; and N. Chomsky, *Language in a psychological setting, Sophia Linguistica* XII (1987). Tokyo: Sophia University. Reprinted by permission of the author.

properties of then-unknown, perhaps still unknown, physical mechanisms. This abstract inquiry served as an essential preliminary and guide for the subsequent inquiry into physical mechanisms. Mentalistic inquiry in the brain sciences is quite similar in approach and character to the abstract inquiry into properties of the chemical elements, and we may expect that this abstract inquiry too will serve as an essential preliminary and guide for the emerging brain sciences today; the logic is quite similar.

This work proceeds from the empirical assumption—which is well-supported—that there is a specific faculty of the mind/brain that is responsible for the use and acquisition of language, a faculty with distinctive characteristics that is apparently unique to the species in essentials and a common endowment of its members, hence a true species property.

These ideas have developed in the context of what some have called "the cognitive revolution" in psychology, and in fact constituted one major factor contributing to these developments. It is important, I think, to understand clearly just what this "revolution" sought to accomplish, why it was undertaken, and how it relates to earlier thinking about these topics. The so-called "cognitive revolution" is concerned with the states of the mind/brain that enter into thought, planning, perception, learning and action. The mind/brain is considered to be an information-processing system, which forms abstract representations and carries out computations that use and modify them. This approach stands in sharp contrast to the study of the shaping and control of behavior that systematically avoided consideration of the states of the mind/brain that enter into behavior, and sought to establish direct relations between stimulus situations, contingencies of reinforcement, and behavior. This behaviorist approach has proven almost entirely barren, in my view, a fact that is not at all surprising since it refuses in principle to consider the major and essential component of all behavior, namely, the states of the mind/brain.

Consider the problem of learning. We have an organism with a mind/brain that is in a certain state or configuration. The organism is presented with certain sensory inputs, leading to a change in the state of the mind/brain. This process is the process of learning, or perhaps more accurately, mental and cognitive growth. Having attained a new state as a result of this process, the organism now carries out certain actions, in part influenced by the state of the mind/brain that has been attained. There is no direct relation between the sensory inputs that led to the change of state of the mind/brain and the actions carried out by the organism, except under highly artificial, uninformative and very marginal conditions.

There is of course a relation of some kind between sensory inputs and behavior; a child who has not been presented with data of Japanese

will not be able to carry out the behavior of speaking Japanese. Presented with appropriate data from Japanese, the child's mind/brain undergoes a significant change; the mind/brain comes to incorporate within itself knowledge of Japanese, which then enables the child to speak and understand Japanese. But there is no direct relation between the data presented to the child and what the child says, and it is hopeless to try to predict what the child will say, even in probabilistic terms, on the basis of the sensory data that led to acquisition of knowledge of Japanese. We can study the process by which the sensory data lead to the change of state of the mind/brain, and we may study at least certain aspects of how this attained knowledge is used. But an effort to study the relation between the sensory data and the actual behavior, avoiding the crucial matter of the nature of the mind/brain and the changes it undergoes, is doomed to triviality and failure, as the history of psychology demonstrates very well. The cognitive revolution was based in part on the recognition of such facts as these, drawing conclusions that really should not be controversial, though they are—a sign of the immaturity of the field, in my view. This change of perspective in the study of psychology, linguistics included, was surely a proper one in essence, and in fact was long overdue.

Not only was this change of perspective overdue, but it also was much less of a revolution than many believed. In fact, without awareness, the new perspective revived ideas that had been developed quite extensively centuries earlier. In particular, seventeenth-century science developed a form of cognitive psychology that was quite rich, and basically, I think, on the right track. Descartes's major scientific contribution, perhaps, was his rejection of the neoscholastic idea that perception is a process in which the form of an object imprints itself somehow on the brain, so that if you see a cube, for example, your brain has the form of a cube imprinted in it in some fashion. In place of this fallacious conception, Descartes proposed a representational theory of mind. He considered the example of a blind man with a stick, who uses the stick to touch in sequence various parts of a physical object before him, let us say a cube. This sequence of tactile inputs leads the blind man construct, in his mind, the image of a cube, but the form of the cube is not imprinted in the mind. Rather, the sequence of tactile inputs leads the mind to construct a mental representation of a cube, using its own resources and its own structural principles. Descartes argued that much the same is true of normal vision. A series of stimuli strike the retina, and the mind then forms ideas that provide a conception of the objects of the external world. The mind then carries out various computational processes, as the person thinks about these objects, including processes that enable the person to carry out certain actions involving them: for example, picking up the cube, rotating it, and so on. This is surely the right general approach. It has been revived in recent psychology and

physiology, and by now something is known about how the process takes place, including even some understanding of the physical mechanisms involved in the coding and representation of stimuli.

Descartes also observed that if a certain figure, say a triangle, is presented to a person, then what the person will perceive is a triangle, though the presented image is certainly not a Euclidean triangle, but rather some far more complex figure. This will be true, he argued, even if the person is a child who has had no previous acquaintance with geometrical figures. In a certain sense the point is obvious, since true geometrical figures do not exist in the natural environment in which we grow and live, but we nevertheless perceive figures as distorted geometrical figures, not as exact instances of whatever they may happen to be. Why does the child perceive the object as a distorted triangle, rather than as the very complex figure that it actually is: with one of the lines slightly curved, with two sides not quite touching, and so on? Descartes' answer was that the Euclidean triangle is produced by the mind on the occasion of this stimulation, because the mechanisms of the mind are based on principles of Euclidean geometry and produce these geometrical figures as exemplars or models for the organization of perception, and for learning, drawing them from its own resources and structural principles.

In contrast, empiricists such as David Hume argued that we simply have no idea of a triangle, or a straight line, since we could not distinguish "perfect images" of such objects from the "defective ones" of the real world. Hume correctly drew the consequences of the empiricist principles that he adopted and developed: in particular, the principle that the mind receives impressions from the outside world and forms associations based upon them, and that this is all there is to the story (apart from the animal instinct underlying induction). But the consequences that Hume correctly drew from these assumptions are certainly false. Contrary to what he asserted, we do, indeed, have a clear concept of a triangle and a straight line, and we perceive objects of the world in terms of these concepts, just as Descartes argued. The conclusion, then, is that the empiricist assumptions are fundamentally wrong, as a matter of empirical fact; the properties of the mind/brain that are involved in determining how we perceive and what we perceive are crucially different from what was postulated in empirical speculation. It seems reasonable to resort to a representational theory of mind of the Cartesian sort, including the concept of the mind as an information-processing system that computes, forms and modifies representations, and we should also adopt something like the Cartesian concept of innate ideas as tendencies and dispositions, biologically determined properties of the mind/brain that provide a framework for the construction of mental representations, a framework that then enters into our perception and action. Ideas of this sort have been revived in the context of the cognitive revolution of the past generation.

Seventeenth-century psychologists, who we call "philosophers," went far beyond these observations. They developed a form of what much later came to be called "Gestalt psychology" as similar ideas were rediscovered during this century. These seventeenth-century thinkers speculated rather plausibly on how we perceive objects around us in terms of structural properties, in terms of our concepts of object and relation, cause and effect, whole and part, symmetry, proportion, the functions served by objects and the characteristic uses to which they are put. We perceive the world around us in this manner, they argued, as a consequence of the organizing activity of the mind, based on its innate structure and the experience that has caused it to assume new and richer forms. "The book of nature is legible only to an intellectual eye," as Ralph Cudworth argued, developing such ideas as these. Again, these speculations seem to be very much on the right track, and the ideas have been rediscovered and developed in contemporary psychology, in part within the context of the cognitive revolution.

The contemporary cognitive revolution has been considerably influenced by modern science, mathematics and technology. The mathematical theory of computation, which developed in the 1920s and 1930s particularly, provided conceptual tools that make it possible to address certain classical problems of representational psychology in a serious way, problems of language in particular. Wilhelm von Humboldt understood, a century and a half ago, that language is a system that makes infinite use of finite means, in his phrase. But he was unable to give a clear account of this correct idea, or to use it as the basis for substantive research into language. The conceptual tools developed in more recent years make it possible for us to study the infinite use of finite means with considerable clarity and understanding. Modern generative grammar, in fact, can be regarded in part as the result of the confluence of the conceptual tools of modern logic and mathematics and the traditional Humboldtian conception, inevitably left vague and unformed. A generative grammar of a language is a formal system that states explicitly what are these finite means available to the mind/brain, which can then make infinite, unbounded use of these means. Unfortunately, the classical ideas concerning language and representational psychology had long been forgotten when the cognitive revolution took place in the 1950s, and the connections I am now discussing were discovered only much later, and are still not widely known.

The development of electronic computers has also influenced the cognitive revolution considerably, primarily in providing useful concepts such as internal representation, modular structure, the software-hardware distinction and the like, and also, in areas such as vision at least, in making it possible to develop explicit models of cognitive processes that can be tested for accuracy and refined. It is worthy of note that much the same was true of the seventeenth-century cognitive revolution. The Cartesians were much impressed with the mechanical

automata then being constructed by skilled craftsmen, which seemed to mimic certain aspects of the behavior of organisms. These automata were a stimulus to their scientific imagination much in the way that modern electronic computers have contributed to the contemporary cognitive revolution.

Some of these seventeenth-century ideas, which are now being rediscovered and developed in quite new ways, have much earlier origins. What is probably the world's first psychological experiment is described in the Platonic dialogues, when Socrates undertakes to demonstrate that a slave boy, who has had no instruction in geometry, nevertheless knows the truths of geometry. Socrates demonstrates this by asking the slave boy a series of questions, providing him with no information but drawing from the inner resources of the slave boy's mind, and in this way Socrates leads the slave boy to the point where he recognizes the truth of theorems of geometry. This experiment was understood, quite plausibly, to show that the slave boy knew geometry without any experience. Indeed, it is difficult to see what other interpretation can be given. The experiment was, presumably, a kind of "thought experiment," but if it were carried out rigorously, as has never been done, the results would probably be more or less as Plato presented them in this literary version of a psychological experiment.

The human mind, in short, somehow incorporates the principles of geometry, and experience only serves to bring them to the point where this innate knowledge can be used. This demonstration also poses a very crucial problem: the problem is to explain how the slave boy can have the knowledge he does have, when he has had no relevant experience from which he could derive this knowledge. Let us refer to this problem as "Plato's problem," returning to it directly.

The rise of generative grammar in the 1950s, a major factor in the cognitive revolution, also resurrected traditional ideas. The Cartesians, in particular, had applied their ideas on the nature of the mind to the study of language, which was commonly viewed as a kind of "mirror of mind." Subsequent study enriched these investigations in quite impressive ways, which we are now only beginning to understand. The cognitive revolution of the 1950s, then, should be understood, I believe, as having recovered independently the insights of earlier years, abandoning the barren dogmas that had impeded understanding of these questions for a very long period; and then applying these classical ideas, now reconstructed in a new framework, in new ways, and developing them along lines that would not have been possible in an earlier period, thanks to new understanding in the sciences, technology and mathematics.

From the point of view adopted in this "second cognitive revolution," the central problems of the study of language are essentially the following four:

The first question, a preliminary to any further inquiry, is this: What is the system of knowledge incorporated in the mind/brain of a person who speaks and understands a particular language? What constitutes the language that the person has mastered and knows? A theory concerned with this topic for a particular language is called "a grammar of that language," or in technical terms, "a generative grammar of the language," where the term "generative grammar" means nothing more than a theory of the language that is fully explicit, so that empirical consequences can be derived in it. Traditional grammars, in contrast, relied crucially on the knowledge of language of the reader of the grammar to fill in the enormous gaps that were left unstudied, and were not even recognized to be gaps; it is surprising, in retrospect, to see how difficult it was to recognize that even the simplest of phenomena pose rather serious problems of explanation. A traditional grammar, then, is not a theory of the language, but is rather a guide that can be followed by a person who already knows the language. Similarly, a pedagogic grammar of Spanish written in English is not a theory of Spanish but rather a guide to Spanish that can be used by a speaker of English who already knows the basic principles of language, though unconsciously, and can therefore make use of the hints and examples in the grammar to draw conclusions about Spanish. A generative grammar, in contrast, seeks to make explicit just what this knowledge is that enables the intelligent reader to make use of a grammar.

To the extent that we can provide at least a partial answer to the first problem, we can turn to a second problem: How is this knowledge of language used in thought or expression of thought, in understanding, in organizing behavior, or in such special uses of language as communication, and so on? Here we have to make a crucial conceptual distinction between (1) the language, a certain cognitive system, a system of knowledge incorporated in the mind/brain and described by the linguist's generative grammar; and (2) various processing systems of the mind/brain that access this knowledge in one or another way, and put it to use.

Still assuming some kind of answer to the problem of characterizing the knowledge attained, we can turn to a third problem: what are the physical mechanisms that exhibit the properties that we discover in the abstract investigation of language and its use; that is, the physical mechanisms of the brain that are involved in the representation of knowledge and in accessing and processing this knowledge? These are pretty much tasks for the future, and they are very difficult ones, primarily, because for very good ethical reasons, we do not permit direct experimentation that might enable scientists to investigate these mechanisms directly. In the case of other systems of the mind/brain, such as the visual system, the investigation of mechanisms has proceeded quite far. The reason is that we allow ourselves, rightly or wrongly, to carry

out direct experimentation with cats, monkeys, and so on. Their visual systems are in many ways like our own, so a good deal can be learned about the physical mechanisms of the human visual system in this way. But it appears that the language faculty is a unique human possession in its essentials, and if we were to discover some other organism that shared this faculty in part, we would probably regard it as quasi-human and refrain from direct experimentation. Consequently, the study of physical mechanisms of the language faculty must be studied in much more indirect ways, either by non-intrusive experiments, or by "nature's experiments," such as injury and pathology. Part of the intellectual fascination of the study of language is that it must proceed in such indirect ways, relying very heavily on the abstract level of inquiry—a difficult and challenging task, but one that can be addressed and has much promise.

The fourth problem is to explain how the knowledge of language and ability to use it are acquired. This problem of acquisition arises both for the language—the cognitive system itself—and for the various processing systems that access the language. I will focus attention here on the first of these questions: on acquisition of language. Plainly, the question can be formulated only to the extent that we have some understanding of what is acquired—of what is a language—though as always, inquiry into the acquisition or use or physical basis of some abstract system can and should provide insight into its nature.

The fourth question is a special case of Plato's problem: How do we come to have such rich and specific knowledge, or such intricate systems of belief and understanding, when the evidence available to us is so meager? That was the problem that rightly troubled Plato, and it should trouble us as well. It is a question that for a long period did not trouble psychologists, linguists, philosophers, and others who thought about the matter, except for a few, who were rather marginal to the main intellectual tradition. This is a sign of the serious intellectual failings of the thought of this era, an interesting topic that I will not pursue here. If a rational Martian scientist were to observe what takes place in a single language community on earth, he would conclude that knowledge of the language that is used is almost entirely inborn. The fact that this is not true, or at least not entirely true, is extremely puzzling, and raises many quite serious problems for psychology and biology, including evolutionary biology.

Recall that Plato had an answer to the problem he posed: we remember the knowledge we have from an earlier existence. This is not a proposal that we would nowadays be inclined to accept in exactly these terms, though we should, in all honesty, be prepared to recognize that it is a far more satisfactory and rational answer than the ones that have been offered in the dominant intellectual traditions of recent centuries, including the Anglo-American empiricist tradition, which simply evaded the problems. To render Plato's answer intelligible, we have to

provide a mechanism by which our knowledge is remembered from an earlier existence. If we are disinclined to accept the immortal soul as the mechanism, we will follow Leibniz in assuming that Plato's answer is on the right track, but must be, in his words, "purged of the error of preexistence." In modern terms, that means reconstructing Platonic "remembrance" in terms of the genetic endowment, which specifies the initial state of the language faculty, much as it determines that we will grow arms not wings, undergo sexual maturation at a certain stage of growth if external conditions such as nutritional level permit this internally directed maturational process to take place, and so on. Nothing is known in detail about the mechanisms in any of these cases, but it is now widely and plausibly assumed that this is the place to look. At least, it is widely assumed for physical growth. The fact that similar evidence does not lead to similar rational conclusions in the case of the mind/brain again reflects the serious intellectual inadequacies of recent thought, which has simply refused to approach problems of the mind/brain by the methods of rational inquiry taken for granted in the physical sciences. This is strikingly true, particularly, of those who falsely believe themselves to be scientific naturalists, and who see themselves as defending science against the obscurantists. Exactly the opposite is true, in my opinion, for the reasons that I have briefly indicated.

Putting aside various dogmas, let us approach questions of mind/brain, including questions of language, in the spirit of the natural sciences. Abstracting away from unknown mechanisms, we assume that the language faculty has an initial state, genetically determined, common to the species apart from gross pathology, and apparently unique to the human species. We know that this initial state can mature to a number of different steady states—the various attainable languages—as conditions of exposure vary. The process of maturation from the initial state to the steady state of mature knowledge is, to some extent, data-driven; exposed to data of English, the mind/brain will incorporate knowledge of English, not Japanese. Furthermore, this process of growth of the language faculty begins remarkably early in life. Recent work indicates that four-day-old infants can already distinguish somehow between the language spoken in their community and other languages, so that the mechanisms of the language faculty begin to operate and to be "tuned" to the external environment very early in life.

It is fairly clear that the process of maturation to the steady state is deterministic. Language learning is not really something that the child does; it is something that happens to the child placed in an appropriate environment, much as the child's body grows and matures in a predetermined way when provided with appropriate nutrition and environmental stimulation. This is not to say that the nature of the environment is irrelevant. The environment determines how the options left unspecified by the initial state of the language faculty are fixed, yielding different languages. In a somewhat similar way, the early

visual environment determines the density of receptors for horizontal and vertical lines. Furthermore, the difference between a rich and stimulating environment and an impoverished environment may be substantial, in language acquisition as in physical growth—or more accurately, as in other aspects of physical growth, the acquisition of language being simply one of these aspects. Capacities that are part of our common human endowment can flourish, or can be restricted and suppressed, depending on the conditions provided for their growth.

The point is probably more general. It is a traditional insight, which merits more attention than it receives, that teaching should not be compared to filling a bottle with water, but rather to helping a flower to grow in its own way. As any good teacher knows, the methods of instruction and the range of material covered are matters of small importance as compared with the success achieved in arousing the natural curiosity of the students and stimulating their interest in exploring on their own. What the student learns passively will be quickly forgotten. What students discover for themselves, when their natural curiosity and creative impulses are aroused, will not only be remembered, but will be the basis for further exploration and inquiry, and perhaps significant intellectual contributions. The same is true in other domains as well. A truly democratic community is one in which the general public has the opportunity for meaningful and constructive participation in the formation of social policy: in their own immediate community, in the workplace, and in the society at large. A society that excludes large areas of crucial decision-making from public control, or a system of governance that merely grants the general public the opportunity to ratify decisions taken by the elite groups that dominate the private society and the state, hardly merits the term "democracy." These too are insights that were alive and vital during the eighteenth century, and have in recent years been largely forgotten or suppressed. The point was made, in another context, by Kant, defending the French Revolution during the period of the Terror against those who argued that the masses of the population "are not ripe for freedom." "If one accepts this proposition," he wrote, "freedom will never be achieved, for one can not arrive at the maturity for freedom without having already acquired it; one must be free to learn how to make use of one's powers freely and usefully . . . one can achieve reason only through one's own experience and one must be free to be able to undertake them. . . . To accept the principle that freedom is worthless for those under one's control and that one has the right to refuse it to them for ever, is an infringement of the rights of God himself, who has created man to be free." Reason, the ability to make use of one's powers freely and usefully, and other human qualities can be achieved only in an environment in which they can flourish. They cannot be taught by coercive means. What is true of physical growth holds quite generally of human maturation and learning.

Returning to the language faculty, learning of language, as noted, is something that happens to the child, without awareness for the most part, just as other processes such as sexual maturation happen to the child. A child does not decide to undergo sexual maturation because it sees others doing so and thinks this would be a good idea, or because it is trained or reinforced. Rather, the process happens in its own inner-directed way. The course of the process, its timing, and its detailed nature are in part influenced by the environment, by nutritional level for example, but the process itself is inner-directed in its essentials. The same appears to be true of language learning, and of other aspects of cognitive growth as well. The term "learning" is, in fact, a very misleading one, and one that is probably best abandoned as a relic of an earlier age, and earlier misunderstandings. Knowledge of language grows in the mind/brain of a child placed in a certain speech community.

Knowledge of language within a speech community is shared to remarkably fine detail, in every aspect of language from pronunciation to interpretation. In each of these aspects, the knowledge attained vastly transcends the evidence available in richness and complexity, and in each of these aspects, the fineness of detail and the precision of knowledge goes well beyond anything that can be explained on any imaginable functional grounds, such as the exigencies of communication. For example, children mimic the sounds of the language around them to a level of precision that is well beyond the capacity of adults to perceive, and in other domains as well, the precision of knowledge and understanding, as well as its scope and richness, are far beyond anything that could be detected in normal human interchange. These properties of normal language can often only be discovered by careful experiment. These are the basic and simplest elements of the problem we face.

We therefore conclude that the initial stage of the language faculty can be regarded as in effect a deterministic input-output system that takes presented data as its input and produces a highly structured cognitive system of a very specific form as its "output"—here the output is internalized, represented in the mind/brain; it is the steady state of knowledge of some particular language. The initial state of the language faculty can be regarded, in essence, as a language-acquisition device; in normal terms, a function that maps presented data into a steady state of knowledge attained. This general conclusion allows many specific variants, to some of which I will briefly return, but it is virtually inconceivable that it is wrong in any fundamental way. There has been much debate over this issue in the literature—more accurately, a one-sided debate in which critics argue that the idea has been refuted, with little response from its defenders. The reason for the lack of response is that the criticism must be based on profound confusion, and inspection of the arguments quickly reveals that this is the case, as it must be, given the nature of the problem.

The theory of the initial state—of the language acquisition device—is sometimes called "universal grammar," adapting a traditional term to a somewhat different conceptual framework. It is commonly assumed that universal grammar, so conceived, determines the class of attainable languages. Let me quote from a recent paper by the two leading researchers in the important new field of mathematical learning theory, a paper on models of language acquisition. They write that universal grammar

> imposes restrictions on a [particular] grammar in such a way that the class of [particular] grammars admissible by the theory includes grammars of all and only natural languages, [where] the natural languages are identified with the languages that can be acquired by normal human infants under casual conditions of access to linguistic data.

The first of these propositions is a definition, and a proper and useful one, so it is not open to challenge: we may define a "natural language" as one that accords with the principles of universal grammar. But the second of these propositions need not be correct. The languages attainable under normal conditions of access are those that fall in the intersection of two sets: (1) the set of natural languages made available by the initial state of the language faculty as characterized by universal grammar, and (2) the set of learnable systems. If universal grammar permits unlearnable languages, as it might, then they simply will not be learned. Learnability, then, is not a requirement that must be met by the language faculty.

Similarly, parsability—that is, the ability of the mind/brain to assign a structural analysis to a sentence—is not a requirement that must be met by a language, contrary to what is often claimed. In fact, we know that the claim is false: every language permits many different categories of expressions that cannot be used or understood readily (or at all), though they are perfectly well-formed, a fact that in no way impedes communication. Furthermore, deviant expressions may be readily parsable, and are often quite properly used. In brief, it is a mistake to think that languages are "designed" for ease of use. In so far as their structure does not conform to functional requirements, their elements are not used.

In the case of learnability, the proposition that the natural languages are learnable may very well be true, but if so, that is not a matter of principle, but rather a surprising empirical discovery about natural language. Recent work in linguistics suggests that it probably is true, again, a surprising and important empirical discovery, to which I will briefly return.

There has been a fair amount of confusion about these matters, in part resulting from misinterpretation of properties of formal systems: for example, the well-known observation that unconstrained transformational grammars can generate all sets that can be specified by finite

means, and results on efficient parsability of context-free languages. In both cases, entirely unwarranted conclusions have been drawn about the nature of language. In fact, no conclusions at all can be drawn with regard to language, language learning, or language use, on the basis of such considerations as these, though other directions of formal inquiry perhaps show more promise of potential empirical significance; for example, some recent work in complexity theory.

When the study of language is approached in the manner I have just outlined, one would expect a close and fruitful interaction between linguistics proper and the investigation of such topics as language processing and acquisition. To some extent this has happened, but less so than might have been hoped. It is useful to reflect a little about why this has been the case. One reason, I think, is the one just mentioned: misinterpretation of results about formal systems has caused considerable confusion. Other problems have arisen from a failure to consider carefully the conceptual relations between language and learnability, and between language and processing. One instructive example is the history of what was called "the derivational theory of complexity," the major paradigm of psycholinguistic research in the early days of the "cognitive revolution." This theory led to an experimental program. The experiments carried out were tests of a theory with two components: (1) assumptions about the rule systems of natural language; (2) assumptions about processing. Some of the experimental results confirmed this combination of theories, others disconfirmed it. But care must be taken to determine just which elements of the combination of theories were confirmed or disconfirmed. In practice, where predictions were disconfirmed, it was concluded that the linguistic component of the amalgam was at fault. While this might be true, and sometimes was as other evidence showed, it was a curious inference, since there was independent evidence supporting the assumptions about language but none whatsoever supporting the assumptions about processing, assumptions that were, furthermore, not particularly plausible except as rough first approximations. Failure to appreciate these facts undermined much subsequent discussion. Similar questions arise with language acquisition, and confirming evidence too, in both areas, is unclear in its import unless the various factors entering into the predictions are properly sorted out.

The history of the derivational theory of complexity illustrates other problems that have impeded useful interaction between linguistics and experimental psychology. Early experimental work was designed to test certain ideas about rule systems on the assumption that processing satisfies the conditions of the derivational theory of complexity. By the time the experimental program had been carried out, with mixed results, the theories of rule systems had changed. Many experimental psychologists found this disconcerting. How can we carry out experi-

mental tests of a theory if it is not stable and is subject to change? These reactions led to a noticeable shift in focus to work in areas that are better insulated from theoretical modification elsewhere.

There are a number of problems with such reactions. One problem is a point of logic: to insulate one's work from theoretical modifications elsewhere is to keep to topics of limited significance, close to the surface of phenomena. If one's work is important enough to have consequences beyond its immediate scope, then it cannot be immune to new understanding outside of this scope. For example, it is likely that results on order of acquisition of function words or on turn-taking in conversation will be immune to discoveries and new understanding elsewhere; the reason is that the implications are very slight. Relevance, after all, is a two-way street. This reaction to the inevitable changes in theoretical assumptions in a discipline that is alive also reflects a far too limited conception of the work of the experimental psychologist, who is perceived as someone who tests ideas developed elsewhere but does not contribute otherwise to their proper formulation. But research into language should obviously be a cooperative enterprise, which can be informed and advanced by use of evidence of many different kinds. There is no privileged sector of this discipline that provides theories, which are tested by others. One sign that the discipline is approaching a higher level of maturity will be that research into language processing and language acquisition will yield conclusions about the structure of language that can be tested by linguists, using the tools of their specific approach to a common network of problems and concerns. The idea that linguistics should be related to psychology as theoretical physics is related to experimental physics is senseless and untenable, and has, I think, been harmful.

Theories of language have indeed undergone significant changes during the period we are now considering—which is to say that the discipline is alive. I think we can identify two major changes of perspective during this period, each with considerable ramifications for the study of language use and acquisition. Let me review these changes briefly, focusing on the three central questions that I mentioned earlier: (1) what is knowledge of language?; (2) how is it acquired?; and (3) how is it used?

Some thirty years ago the standard answers to these questions would have been something like this.

1. What is knowledge of language? Answer: it is a system of habits, dispositions and abilities. This answer, incidentally, is still widely held, notably by philosophers influenced by Wittgenstein and Quine.
2. How is language acquired? Answer: by conditioning, training, habit-formation or "general learning mechanisms" such as induction.
3. How is language used? Answer: language use is the exercise of an ability, like any skill; say, bicycle-riding. New forms are produced or

understood "by analogy" to old ones. In fact, the problem posed by production of new forms, the normal situation in language use, was barely noticed. This is quite a remarkable fact, first, because the point is obvious, and second, because it was a major preoccupation of the linguistics of the first cognitive revolution of the seventeenth century. Here we have a striking example of how ideology displaced the most obvious of phenomena from inquiry.

Attention to the simplest phenomena suffices to show that these ideas cannot be even close to the truth of the matter, and must simply be abandoned. Let me illustrate with a very simple example. Imagine a child learning English who comes to understand the sentence *John ate an apple*. The child then knows that the word *eat* takes two semantic roles, that of the subject (the agent of the action) and that of object (the recipient of the action); it is a typical transitive verb. Suppose that the child now hears the reduced sentence *John ate*, in which the object is missing. Since the verb is transitive, requiring an object, the child will understand the sentence to mean, roughly, "John ate something or other." So far everything is fairly straightforward if we assume the simple principle that when a semantically required element is missing, the mind interprets it to be a kind of "empty pronoun" meaning: something or other. Perhaps an empiricist linguist might be willing to suppose that this principle is available as an innate element of the language faculty.

Consider now a very simple but slightly more complex sentence. Suppose the child comes to understand such sentences as *John is too clever to catch Bill*. Here the verb *catch* also requires a subject and an object, but the subject is missing in this sentence. It therefore has to be supplied by the mind, in the manner of the object of *ate* in *John ate*. By the principle just assumed to account for *John ate*, the sentence should mean: John is so clever that someone or other will not catch Bill. That is a fine meaning, but it is not the meaning of *John is too clever to catch Bill*. Rather, the sentence means: John is so clever that he, John, will not catch Bill. The mind does not use the empty pronoun principle, but rather takes the subject of *catch* to be the same as the subject of *is clever*. Since this is known without instruction or evidence, we must attribute to the mind still a second principle, let us call it the principle of subject control: the missing subject of the embedded clause is understood to be the same as the subject of the main clause. Our assumptions about the innate resources of the mind must therefore be enriched.

Let us carry the discussion a step further. Suppose we delete *Bill* from the sentence *John is too clever to catch Bill*, so that we have *John is too clever to catch*. By the empty pronoun principle and the subject control principle, the sentence should mean: John is so clever that he, John, will not catch someone or other. But the child knows that it does not mean that at all; rather, it means that John is so clever that someone or

other will not catch him, John. The child interprets the sentence by some other principle, call it the inversion principle, which tells us that the object of the embedded sentence is understood to be the same as the subject of the main verb, and the subject of the embedded sentence is an empty pronoun referring to someone or other.

We now have to attribute to the mind/brain three principles: the empty pronoun principle, the subject principle, and the inversion principle. Furthermore, some overarching principle of the mind/brain determines when these principles of interpretation are applied.

Turning to slightly more complicated examples, the mysteries deepen. Consider the sentence *John is too clever to expect anyone to catch*. English speakers at first may find this sentence a bit puzzling, but "on reflection" (whatever that involves), they understand it to mean that John is so clever that someone doesn't expect anyone to catch John; that is, it is interpreted by means of the empty pronoun principle and the inversion principle. But now compare this sentence with another that is roughly comparable in complexity: *John is too clever to meet anyone who caught*. Here all principles fail; the sentence is complete gibberish. We can parse the sentence with no difficulty; it just doesn't mean anything sensible. In particular, it is not understood "by analogy" to mean that John is so clever that no one met anyone who caught him, John.

Notice that none of this is the result of training, or even experience. These facts are known without training, without correction of error, without relevant experience, and are known the same way by every speaker of English—and in analogous constructions, other languages. Hence all of this must somehow derive from the inner resources of the mind/brain, from the genetically determined constitution of the language faculty. Clearly the answer cannot be that these resources include the empty pronoun principle, the subject principle, the inversion principle, some principle that determines how they operate, and a principle blocking the "analogy" in the last example. Rather, we would like to show that the observed facts follow from some deeper principles of the language faculty. This is a typical problem of science, and one that has, in fact, been rather successfully addressed in recent work. But the point here is that the facts show rather clearly that the standard answers to our questions that I have just mentioned cannot be on the right track.

Notice again that the concept of "analogy" does no work at all. By analogy to *John ate*, the sentence *John is too clever to catch* should mean "John is too clever to catch someone or other," but it does not. Notice also that such examples refute the conception of knowledge of language as a skill or ability. The child does not fail to provide the analogous interpretation because of a failure of ability—because it is too weak, or needs more practice. Rather, the computational system of the mind/brain is designed to force certain interpretations for linguistic expressions. To put the matter in the context of the theory of knowledge, our

knowledge that expression such-and-such means so-and-so is not justified or grounded in experience in any useful sense of these terms, is not based on good reasons or reliable procedures, is not derived by induction or any other general method. Since these are examples of ordinary propositional knowledge, knowledge that so-and-so, the standard paradigms of epistemology and fixation of belief cannot be correct, and investigation of further examples and other cognitive systems reveals exactly the same thing, so I believe.

I think that these are all important facts, insufficiently appreciated, with quite considerable import. We discover facts of this sort wherever we look, if we are not blinded or misled by dogma.

One notable feature of the widely held conceptions of knowledge and language in terms of ability, skill, habit, general learning mechanisms and analogy, is that they were entirely unproductive and without empirical consequences. One can hardly point to a single empirical result of the slightest significance that derived from these conceptions. The psychology of language of the time was almost completely barren. There was an empirical discipline, namely structural linguistics, which did profess these doctrines and did achieve empirical results and some theoretical understanding. But a closer look will show that in practice, research departed from the professed ideology at every crucial point. The general conceptual framework limited and impoverished the discipline, barring natural lines of inquiry, but otherwise was simply professed and abandoned in practice, though it did, I believe, have a serious and generally harmful impact on applied disciplines such as language teaching.

Recognition of the complete inadequacy of these conceptions led to the first major conceptual change, which was, in many respects, a return to traditional ideas and concerns that had been dismissed or forgotten during the long period when empiricist and behaviorist doctrines prevailed. This shift of focus provided a new set of answers to the central question:

1. What is knowledge of language? Answer: language is a computational system, a rule system of some sort. Knowledge of language is knowledge of this rule system.

2. How is language acquired? Answer: the initial state of the language faculty determines possible rules and modes of interaction. Language is acquired by a process of selection of a rule system of an appropriate sort on the basis of direct evidence. Experience yields an inventory of rules, through the language-acquisition device of the language faculty.

3. How is language used? Answer: the use of language is rule-governed behavior. Rules form mental representations, which enter into our speaking and understanding. A sentence is parsed and understood

by a systematic search through the rule system of the language in question.

The new set of answers constitutes a major component of the "cognitive revolution."

This was a significant shift of point of view: from behavior and its products to the system of knowledge represented in the mind/brain that underlies behavior. Behavior is not the focus of inquiry; rather, it simply provides one source of evidence for the internal systems of the mind/brain that are what we are trying to discover—the system that constitutes a particular language and that determines the form, structural properties and meaning of expressions, and more deeply, the innate structure of the language faculty. As I mentioned earlier, this shift towards an avowed mentalism is also a shift towards assimilating the study of language to the natural sciences, and opens up the possibility of a serious investigation of physical mechanisms.

This shift of focus was extremely productive. It led to a rapid decrease in the range of empirical phenomena that were brought under investigation, with many new empirical discoveries, such as those just illustrated, including very simple facts that had never been noticed. It also led to some degree of success in providing explanations for these facts. But serious difficulties arise at once. Basically, these relate to Plato's problem, the problem of acquisition of language. In essence, the problem is that there are too many possible rule systems. Therefore it is hard to explain how children unerringly select one such system rather than another. Furthermore, children seem to select very complex rule systems and systematically to avoid much simpler ones, a conclusion that makes no sense.

These problems set the research agenda since about 1960, within the framework I am considering here. I will not review the steps that were taken, but rather will turn to the result. In the past several years, a new and very different conception of language has emerged, which yields new answers to our three questions. The initial state of the language faculty consists of a collection of subsystems, or *modules* as they are sometimes called, each of which is based on certain general principles. Many of these principles admit of a certain limited possibility of variation. We may think of the system as a complex network associated with a switch box that contains a finite number of switches. The network is invariant, but each switch can be in one of several positions, perhaps two: on or off. Unless the switches are set, nothing happens. But when the switches are set in one of the permissible ways, the system functions, yielding the entire infinite array of interpretations for linguistic expressions. A slight change in switch settings can yield complex and varied phenomenal consequences as its effects filter through the network. There are no rules at all, hence no necessity to learn rules. For example, the possible phrase structures of a language are fixed by

general principles and are invariant among languages, but there are some switches to be set. One has to do with order of elements. In English, for example, nouns, verbs, adjectives and prepositions precede their objects: in Japanese, the comparable elements follow their objects. English is what is called a "head-first" language, Japanese a "head-last" language. These facts can be determined from very simple sentences; for example, the sentences "John ate an apple" (in English) or "John an apple ate" (in Japanese). To acquire a language, the child's mind must determine how the switches are set, and simple data must suffice to determine the switch settings, as in this case. The theory of language use also undergoes corresponding modifications, which I cannot explore here.

This second conceptual change gives a very different conception of language and knowledge. To mention one example, notice that from the point of view of rule systems, there are an infinite number of languages, since there are infinitely many rule systems of the permissible form. But from the network-switch point of view, there are only finitely many languages, one for each arrangement of switch settings. Since each of the switch settings can be determined from simple data, each of these finitely many languages is learnable. Hence the general principle of learnability theory discussed earlier is in fact true: each natural language is learnable—though it is far from true that the learnable systems are all natural languages. As I mentioned, this is an empirical result, and a very surprising one, not a matter of principle. There is, incidentally, some intriguing work in mathematical learning theory which suggests that language acquisition is possible in principle under plausible conditions only if the set of natural languages is indeed "finite" (in a special sense).

This second conceptual change has, once again, led to a great increase in the range of empirical materials discovered and subjected to serious inquiry within generative grammar, now from a much wider range of languages.

Assuming that this change is pointing in the right direction, what are the consequences for the study of language acquisition? The problem will be to determine how the switches are set and to discover the principles of learning, or maturation, or whatever is responsible for carrying out the transition from the initial state of the language faculty to the steady state of adult competence; that is, for setting the switches of the language faculty. Recall that two factors enter into language acquisition: the nature of the language faculty, and the principles of learning theory or more properly growth theory, and any evidence about language acquisition must be assessed carefully to determine how it bears on one or the other of these two interacting factors. How can we proceed in studying this question?

Notice that the problem of assessment of evidence and explanation would plainly be simplified if one or the other of these two compo-

nents—universal grammar or growth theory—does not exist. Each of these positions has been maintained, the first one quite vigorously, the second as a tentative working hypothesis.

Denial of the existence of universal grammar—that is, of the language faculty as an identifiable system of the human mind/brain—is implicit in the empiricist program and in some recent claims about mechanisms of "general intelligence" or "connectionism" or theory formation, mechanisms that are allegedly applied to yield our linguistic abilities and other intellectual achievements in an undifferentiated way. There has been no attempt to formulate these alleged mechanisms that seems to offer any real promise. The clearer formulations have been quickly refuted, in some cases refuted in principle, and for reasons that should be familiar, the prospects for this program seem very dim. Since there is nothing substantive to discuss, I will disregard this possibility and proceed to the second possibility: that growth theory is negligible or non-existent, so that language-acquisition simply involves switch setting on the basis of presented data, such as the sentences "John ate an apple" and "John an apple ate." Let us call this the "no-growth theory" of language acquisition.

Obviously, this cannot be literally true. During the first few months or perhaps weeks of life, an infant probably is exposed to enough linguistic data to set most switches, but plainly it has not done so. In fact, the process extends over quite a few years. So to maintain the no-growth theory we would have to argue that some independent and extrinsic change in cognitive capacities, say in memory or attention, accounts for the observed stages of growth.

Such ideas have been advanced with regard to stages of cognitive development in the sense of Jean Piaget, and also with respect to the stages of language growth. For example, it has been observed that the transition from so-called "telegraphic speech," lacking function words, to normal speech is quite rapid, and includes a number of different systems: questions, negations, tag questions, etc. Furthermore, in the telegraphic speech stage, children understand normal speech better than their own telegraphic speech, and if function words are introduced randomly, the results are unintelligible. This suggests that the children knew the facts of normal speech all along, and were using telegraphic speech because of some limitation of attention and memory. When this limitation is overcome in the course of normal growth and maturation, their already acquired knowledge of language can be manifested. But there are some serious problems in assuming this idea in other cases of regular stages of development: for example, the shift from semantic to syntactic categories, the use of color words, the appearance of a true verbal passive construction and other more complex structures, the emergence of semantic properties of control, and so on. Prima facie, it seems hard to explain these transitions without appeal to maturational processes that bring principles of universal grammar into operation on

some regular schedule in a manner to be described and accounted for in a genetic theory. Of course, what is prima facie plausible is not necessarily correct, but the questions that arise are clear enough, and it is an important task to address them, as many investigators are now doing in important recent work.

There is, on the one hand, work by Yukio Otsu, Stephen Crain, and others that seems to show that principles of universal grammar are available as soon as constructions are used in which they would be manifested, and the delay in use of these constructions might be explained in terms of inherent complexity, hence extrinsic factors such as memory.

To take one complex example of much general interest, consider recent work of Nina Hyams on the null subject property that distinguishes languages like French and English, in which subjects must be overtly expressed, from languages such as Italian and Spanish, in which the subject may be suppressed in the phonetic output. Hyam's work indicates that at an early stage, all children treat their language as if it were a null subject language. The switch, she suggests, has what is called an "unmarked setting," or in the more usual terminology, the null subject parameter has an "unmarked value," a value selected in the absence of data, and this value provides a null subject language. Italian-speaking children maintain the unmarked value, while English-speaking children later change to the marked value of the parameter, setting the switch differently. The question then is: What triggers the change? There is good evidence that positive evidence suffices for language acquisition; that is, correction of error is unnecessary and probably largely irrelevant when it occurs. Assuming so, the answer to the question cannot be that the English-speaking children are explicitly corrected. Nor can the answer be that they never hear sentences without subjects, since they hear no evidence for most of what they know. Assuming a no-growth theory, Hyams suggests that the change is triggered by the presence of overt expletives in English, such elements as *there* in "there is a man in the room," elements that are semantically empty but must be present to satisfy some syntactic principle. The assumption is that universal grammar contains a principle implying that if a language has overt expletives, then it is not a null subject language. This is, incidentally, an example of a hypothesis about universal grammar deriving from language acquisition studies that might be tested by linguists, rather than the converse, as in the usual practice. It cannot be quite correct as it stands, but something similar might be true.

But now we have to ask why the English-speaking children delay in using this evidence. A possible answer (though not the one Hyams proposes) might be that extrinsic conditions of memory and attention render these expletives inaccessible at an early stage.

Pursuing a similar idea, Luigi Rizzi suggests that contrary to Hyams's initial conclusion, the unmarked value for the parameter is: overt sub-

ject. English-speaking children appear to violate this principle at an early stage, but only because extrinsic considerations suppress the production of such elements as unstressed subject pronouns. Italian-speaking children then select the marked value of the parameter on the basis of direct evidence of subjectless sentences.

A third approach is to reject the no-growth theory and to suppose that the null subject parameter only becomes available at a certain stage of maturation, and is set at the marked null subject value only if direct evidence of subjectless sentences is presented. At the moment, the question remains open, and these possibilities do not exhaust the options (for example, the null subject parameter might be further differentiated, or cast in different terms).

Notice that further clarification of these issues might well contribute to our knowledge of the principles and parameters of universal grammar—of the nature of the network and the switches—on the basis of evidence from language acquisition, as we should anticipate as the discipline progresses.

Consider a second example. Sascha Felix argues against the no-growth theory on the basis of evidence about use of negatives in several languages. Apparently, at the earliest stage, children use sentence-external negation, as in "not John likes milk." This fact (which, not surprisingly, is itself controversial) already raises problems for a no-growth theory, since natural languages rarely if ever exhibit sentence-external negation. At a later stage, the child shifts to sentence-internal negation, as in "John no likes milk," which is also inconsistent with the evidence from the adult language. Later, the correct form "John doesn't like milk" emerges. Felix points out that stage I, with sentence-external negation, is consistent with Dan Slobin's principle that the learner avoids interruption of linguistic units, and hence might be taken to support this principle. But he notes that that leaves unresolved the question why this principle becomes inoperative at stage II, and is even more radically abandoned at stage III. A maturational theory seems a possible candidate for an explanation. Again, further research should contribute to clarifying both the principles of language growth, if they exist, and the actual principles and parameters of universal grammar.

Consider finally a more complex example studied in some detail by Hagit Borer and Kenneth Wexler. They argue that the results in many languages on acquisition of passives can be explained by a maturational theory, which provides a more sophisticated version of the idea that transformations are acquired step-by-step during language acquisition. Their theory postulates that until a certain stage of development, phrases can only be interpreted in a canonical position in which semantic roles are assigned by principles of universal grammar, thus the position of abstract underlying deep structures, in effect. At this stage, a sentence such as "John was killed" is simply uninterpretable, since

John is displaced from its canonical position as object of *kill*. Apparent passive forms at this stage, they argue, are in fact adjectives, as in "the door is closed." Later, a device becomes available, through maturation, by which displaced elements can be interpreted through a so-called *chain* formed by a transformation, which links the displaced element to an empty *trace* in the canonical position. Such chains must then meet various conditions of universal grammar, which account for the possibilities of displacement. They argue that the range of available evidence about acquisition of passives can be largely explained on the basis of this assumption: that chains become available at a certain stage of maturation. Again, there are numerous empirical problems and consequences to be explored, and the results should bear directly on the principles of universal grammar as well as growth theory.

If Borer and Wexler are right, one might be tempted to explore a famous suggestion by Roman Jakobson that language acquisition and language loss in aphasia are mirror images: the earlier some items and structures are acquired in language learning, the later they are lost under brain injury. It would then follow that in some kinds of aphasia, we should find that chains are lost while other aspects of phrase structure remain. Evidence to this effect has in fact been presented by Yosef Grodzinsky. This again suggests what might prove to be an intriguing line of inquiry.

These examples barely scratch the surface. A wide range of intriguing questions arise at once if we think through the implications of the principles-and-parameters conception of universal grammar in terms of an invariant network and an associated set of switches, and if we ask how this conception might relate to possible principles of maturation involved in language growth, along with extrinsic factors in cognitive development. I have not had time to consider the question of language processing, but here too the questions look quite different when approached in these terms. And within the study of language proper, many new and exciting questions enter into the research agenda. If the principles-and-parameters approach is correct, it should be possible literally to deduce the properties of each natural language by setting the switches in one or another way and computing the consequences. Typological difference should be a matter of differences in switch-settings. Language change should be the result of a change in such a setting; note that a small change might yield a substantial phenomenal difference as its effects filter through the fixed network of modular principles. These are all questions that are now being addressed, in some cases with some success, in other cases with failures that are highly suggestive in opening up new lines of inquiry. Furthermore, the class of typologically different languages that have come under investigation, and that seem to be amenable to a coherent and uniform approach, has vastly extended, again, a promising sign.

There are, it seems, real grounds for considerable optimism about the prospects that lie ahead, not only for the study of language proper, but also for the study of cognitive systems of the mind/brain of which language is a fundamental and essential component, in the human species.