

Language and Problems of Knowledge

Noam Chomsky

RESUMEN

Cualquier enfoque serio del estudio del lenguaje parte del uso de sentido común, reemplazándolo por algún concepto técnico. La práctica general ha sido definir “lenguaje” como extensional y externalizado (lenguaje-E). Sin embargo, tal concepto de lenguaje y sus variantes plantea numerosos problemas, y en este artículo se argumenta que el lenguaje-E es un término artificial sin estatus alguno en una teoría del lenguaje. Una táctica mejor sería considerar el “lenguaje” como intensional e internalizado (lenguaje-I). Teniendo presente esta concepción, el enfoque del lenguaje como conocimiento de una capacidad – algo muy común entre los filósofos influenciados por Wittgenstein– se rechaza como pura fantasía. Por otra parte, el abandono de las reglas en favor de principios generales hace que la cuestión de por qué elegimos unas reglas en vez de otras, simplemente no se plantea. Todo esto parece forzarnos a abandonar muchas doctrinas comúnmente aceptadas sobre el lenguaje y el conocimiento. Hay una estructura innata que determina el armazón dentro del cual se desarrollan el pensamiento y el lenguaje, llegando hasta detalles muy precisos e intrincados. El lenguaje y el pensamiento se despiertan en la mente y siguen un curso en gran medida predeterminado, del mismo modo que muchas otras propiedades biológicas.

ABSTRACT

Every serious approach to the study of language departs from common sense usage, replacing it by some technical concept. The general practice has been to define “language” as extensional and externalized (E-language). However, such a concept of language and its variants raises numerous problems and it is argued in this paper that E-language is an artifact with no status in a theory of language. A better move would be to consider “language” as intensional and internalized (I-language). Bearing this conception in mind, the approach to the knowledge of language as knowledge of an ability — very common among those philosophers influenced by Wittgenstein — is examined and rejected as a philosopher’s pipe dream. At the same time, as a result of the abandonment of rules in favour of general principles, the question of why we have one choice of rules rather than another simply does not arise. All this seems to force us to abandon many commonly accepted doctrines about language and knowledge. There is an innate structure that determines the framework within which thought and language develop, down to quite precise and intricate details. Language and thought are awakened in the mind and follow a largely predetermined course, much like other biological properties.

Before entering into the question of language and problems of knowledge, it may be useful to clarify some terminological and conceptual issues concerning the concepts “language” and “knowledge” which, I think, have tended to obscure understanding and to engender pointless controversy.

To begin with, what do we mean by “language”? There is an intuitive common-sense concept that serves well enough for ordinary life, but it is a familiar observation that every serious approach to the study of language departs from it quite sharply. It is doubtful that the common sense concept is even coherent, nor would it matter for ordinary purposes if it were not. It is, in the first place, an obscure sociopolitical concept, having to do with colors on maps and the like, and a concept with equally obscure normative and teleological elements, a fact that becomes clear when we ask what language a child of five, or a foreigner learning English, is speaking — surely not my language, nor any other language, in ordinary usage. Rather, we say that the child and foreigner are “on their way” to learning English, and the child will “get there”, though the foreigner probably will not, except partially. But if all adults were to die from some sudden disease, and children of five or under were to survive, whatever it is that they were speaking would become a typical human language, though one that we say does not now exist. Ordinary usage breaks down at this point, not surprisingly: its concepts are not designed for inquiry into the nature of language.

Or consider the question of what are called “errors”. Many, perhaps most speakers of what we call “English” believe that the word “livid”, which they have learned from the phrase “livid with rage”, means “red” or “flushed”. The dictionary tells us that it means “pale”. In ordinary usage, we say that the speakers are wrong about the meaning of this word of their language, and we would say this even if 95%, or perhaps 100% of them made this “error”. On the other hand, if dictionaries and other normative documents were destroyed with all memory of them, “livid” would then mean “flushed” in the new language. Whatever all this might mean, it plainly has nothing to do with an eventual science of language, but involves other notions having to do with authority, class structure and the like. We understand this easily enough in connection with pronunciation; thus to say that the pronunciation of one dialect is “right” and that of another is “wrong” makes as much sense as saying that it is “right” to talk English and “wrong” to talk Spanish. Such judgments, whatever their status, plainly have nothing to do with the study of language and mind, or human biology. The questions of meaning seem to we quite parallel in status.

Note that a person can be mistaken about his or her own language. Thus if “livid” in fact means “flushed” in my current language, and I tell you that it means “pale” in my language, then I am wrong, just as I would be wrong if I told you, perhaps in honest error, that in my language “whom” is

always used for a direct object, not “who” or if I were to deny some feature of the urban dialect that I speak natively. Judgments about oneself are as fallible as any others, but that is not what is at issue here.

All of this is, or should be, commonplace. Correspondingly, every serious approach to the study of language departs from the common sense usage, replacing it by some technical concept. The choices have generally been questionable ones. The general practice has been to define “language” as what I have called elsewhere “E-language”, where “E” is intended to suggest “extensional” and “externalized”. The definition is “extensional” in that it takes language to be a set of objects of some kind, and it is “externalized” in the sense that language, so defined, is external to the mind/brain. Thus a set, however chosen, is plainly external to the mind/brain.

As a side comment, let me say that I will use mentalistic terminology freely, but without any dubious metaphysical burden; as I will use the terms, talk about mind is simply talk about the brain at some level of abstraction that we believe to be appropriate for understanding crucial and essential properties of neural systems, on a par with discussion in 19th-century chemistry of valence, benzene rings, elements, and the like, abstract entities of some sort that one hoped would be related, ultimately, to then-unknown physical entities.

A typical formulation of a notion of E-language is the definition of “language” by the distinguished American linguist Leonard Bloomfield as “the totality of utterances that can be made in a speech community”, the latter another abstract entity, assumed to be homogeneous¹. Another approach, based ultimately on Aristotle’s conception of language as a relation of sound and meaning, is to define “language” as a set of pairs (s, m) , where s is a sentence or utterance, and m is a meaning, perhaps represented as some kind of set-theoretical object in a system of possible worlds, a proposal developed by the philosopher David Lewis among others. There are other similar proposals.

Under any of these proposals, a grammar will be a formal system of some kind that enumerates or “generates” the set chosen to be “the language”, clearly an infinite set for which we seek a finite representation.

The concept “E-language” and its variants raise numerous questions. In the first place, the set is ill-defined, non simply in the sense that it may be vague, with indeterminate boundaries, but in a deeper sense. Consider what are sometimes called “semi-grammatical sentences”, such as “the child seems sleeping”. Is this in the language or outside it? Either answer is unacceptable. The sentence clearly has a definite meaning. An English speaker interprets it in a definite way, quite differently from the interpretation that would be given by a speaker of Japanese. Hence it cannot simply be excluded from the set “E-English”, though it is plainly not well-formed. But speakers of English and Japanese will also differ in how they interpret some sentence of Hindi — or for that matter how they will interpret a wide variety of noises — so then

all languages and a vast range of other sounds also fall within English, a conclusion that makes no sense. It is doubtful that there is any coherent solution to this range of problems. The fact is that a speaker of English, Japanese, or whatever, has developed a system of knowledge that assigns a certain status to a vast range of physical events, and no concept of E-language, nor any construct developed from it, is likely to be able to do justice to this essential fact.

A second problem has to do with choice of grammar. Evidently, for any set there are many grammars that will enumerate it. Hence it has commonly been argued, most notably by W. v. Quine, that choice of grammar is a matter of convenience, not truth, like the choice of a “grammar” for the well-formed sentences of arithmetic in some notation. But now we face real questions about the subject matter of the study of language. Clearly, there is some fact about the mind/brain that differentiates speakers of English from speakers of Japanese, and there is a truth about this matter, which is ultimately a question of biology. But sets are not in the mind/brain, and grammars can be chosen freely as long as they enumerate the E-language, so the study of E-language, however construed, does not seem to bear on the truth about speakers of English and Japanese; it is not, even in principle, part of the natural sciences, and one might argue that it is a pointless pursuit. Many philosophers — W. V. Quine, David Lewis, and others — have concluded that linguists must be in error when they hold that they are concerned with truths about the mind/brain, though clearly there are such truths about language for someone to be concerned with; they also hold that puzzling philosophical problems are raised by the claim that grammars are “internally represented” in some manner. Others (Jerrold Katz, Scott Soames, and others) have held that linguistics is concerned with some Platonic object that we may call “P-language”, and that P-English is what it is independently of what may be true about the psychological states or brains of speakers of English. One can see how these conclusions might be reached by someone who begins by construing language to be a variety of E-language.

There is little point arguing about how to define the term “linguistics”, but it is plain that there is an area of investigation, let us call it “C-linguistics” (cognitive linguistics), which is concerned with the truth about the mind/brains of the people who speak C-English and C-Japanese, suitably idealized. This subject belongs strictly within the natural sciences in principle, and its links to the main body of the natural sciences will become more explicit as the neural mechanisms responsible for the structures and principles discovered in the study of mind come to be understood. As I noted earlier, the status of this study of language and mind is similar to that of 19th century chemistry or pre-DNA genetics; one might argue that it is similar to the natural sciences at every stage of their development. In any event, C-linguistics raises no philosophical problems that do not arise for scientific in-

quiry quite generally. It raises numerous problems of fact and interpretation, but of a kind familiar in empirical inquiry.

The status of P-linguistics, or of the study of E-language generally, is quite different. Thus the advocates of P-linguistics have to demonstrate that in addition to the real entities C-English, C-Japanese, etc., and the real mind/brains of their speakers, there are other Platonic objects that they choose to delineate somehow and study. Whatever the merits of this claim, we may simply put the matter aside, noting that people may study whatever abstract object they construct. This still leaves the apparent problem noted by Quine, Lewis and others who argue that it is “folly” to claim that one of a set of “extensionally equivalent systems of grammar” that enumerate the same E-language is correctly attributed to the speaker-hearer as a property physically encoded in some manner, whereas another one merely happens to enumerate the E-language but is not a correct account of the speaker’s mind/brain and system of knowledge. Plainly this conclusion cannot be correct, given that as they agree, there is surely some truth about the mind/brain and the system of knowledge represented in it, so some error must have crept in along the way.

Note that the question is not one of metaphysical realism, or of choice of theory in science. Take whatever view one wants on these matters, and it is still alleged that some further philosophical problem, or “folly”, arises in the case of attribution of one grammar but not another extensionally equivalent one to a speaker-hearer, a conclusion that is transparently in error, but seems to be as well-founded as the correct conclusion that there is no “true” grammar of arithmetic. So we seem to be left with a puzzle. .

A third class of problems that arise from the study of E-language has to do with the properties of these sets. Sets have formal properties, so it seems to be meaningful to ask, whether human E-languages have certain formal properties: are they context-free, or recursive, or denumerable? All of these choices have been affirmed, and denied, but the point is that the questions are taken seriously, though in my view they are quite meaningless and empty. The answers are also thought to have some crucial bearing on questions of parsing and learnability, again wrongly.

All of this is, in my view, quite confused and pointless, because the notion of E-language is an artifact, with no status in an eventual science of language. E-languages can be selected one way or another, or perhaps better, not at all, since there appears to be no coherent choice and the concept appears to be useless for any empirical inquiry. In particular, it is quite mistaken to hold, as many do, that an E-language is somehow “given”, and that there is no particular problem in making sense of the idea that a person uses a particular E-language, but that in contrast there are serious problems if not pure folly in the contention that a particular “grammar” for that E-language, but

not some other one, is in fact used by the speaker. Clearly infinite sets are not “given”. What is given to the child is some finite array of data, on the basis of which the child’s mind develops some system of knowledge X , where X determines the status of arbitrary physical objects, assigning to some of them a phonetic form and meaning. With a different finite array of data — from Japanese rather than English, for example — the system of knowledge attained will differ, and the question of what the systems in the mind/brain really are is as meaningful as any other question of science. As for the E-language, it does raise innumerable problems, probably unanswerable ones, since whatever it is, if anything, it is more remote from mechanisms and at a higher level of abstraction than the internally represented system of knowledge, the “correct grammar” that is alleged to raise such difficulties.

The source of all of these problems resides in an inappropriate choice of the basic concept of the study of language, namely “language”. The only relevant notion that has a real status is what is usually called “grammar”. Here again we find an unfortunate terminological decision, which has undoubtedly been misleading. Guided by the misleading and inappropriate analogy to formal languages, I and others have used the term “language” to refer to some kind of E-language, and have used the term “grammar” with systematic ambiguity — a fact that has always been spelled out clearly, but has nevertheless caused confusion: the term “grammar” has been used to refer to the linguist’s theory, or to the subject matter of that theory. A better usage would be to restrict the term “grammar” to the theory of the language, and to understand the language as what we may call “I-language”, where “I” is to suggest “intensional” and “internalized”. The I-language is what the grammar purports to describe: a system represented in the mind/brain, ultimately in physical mechanisms that are now largely unknown, and is in this sense *internalized*; a system that is *intensional* in that it may be regarded as a specific function considered in intension — that is, a specific characterization of a function — which assigns a status to a vast range of physical events, including the utterance “John seems to be sleeping”, the utterance “John seems sleeping”, a sentence of Hindi, and probably the squeaking of a door, if we could do careful enough experiments to show how speakers of English and Japanese might differ in the way they “hear” this noise.

As contrasted with E-language, however construed, I-languages are real entities, as real as chemical compounds. They are in the mind, ultimately the brain. They are what they are, and it is a problem of science to discover the true account of what they are, the grammar for the speaker in question. The story presented by many philosophers is entirely backwards. It is the E-language, not the I-language (the “grammar”, in one of the two senses in which this systematically ambiguous phrase has been used), that poses philosophical problems, which are probably not worth trying to solve, since the

concept is of no interest and has no status. It may, indeed, be pure “folly” to construct and discuss it, to ask what formal properties E-languages have, and so on. I suspect it is. In particular, the analogy to formal systems of arithmetic and so on is largely worthless, and should be discarded, though other analogies to arithmetic and logic, as systems of mentally represented knowledge, are quite definitely worth pursuing, and raise quite interesting questions, yet to be seriously explored. The debates of the past generation about these matters seem to me a classic example of the philosophical errors that arise from accepting a useless, perhaps quite senseless concept, and assuming falsely that it is the one that corresponds, or most properly replaces, some concept of ordinary language, a source of philosophical error that was clearly exposed in the 18th century, if not earlier, and has more recently been brought to general attention by Wittgenstein.

Let us now use the term “language” to refer to I-language, and the term “grammar” to refer to the theory of an I-language. What about the term “universal grammar”, recently resurrected and given a sense that is similar to the traditional one, but not identical, since the entire framework of thinking has been radically modified? The term “universal grammar” has also been used with systematic ambiguity, to refer to the linguist’s theory and to its subject matter. In keeping with our effort to select terms so as to avoid pointless confusion, let us use the term “universal grammar” to refer to the linguist’s theory only. The topic of universal grammar is, then, the system of principles that specify what it is to be a human language. This system of principles is a component of the mind/brain prior to the acquisition of a particular language. It is plausible to suppose that this system constitutes the initial state of the language faculty, considered to be a subsystem of the mind/brain.

This initial state, call it S_0 , is apparently a common human possession to a very close approximation, and also appears to be unique to humans, hence a true species property. It is what it is, and theories concerning it are true or false. Our goal is to discover the true theory of universal grammar, which will deal with the factors that make it possible to acquire a particular I-language and that determine the class of humanly attainable I-languages and their properties. Looked at from a certain point of view, universal grammar describes a “language acquisition device”, a system that maps data into language (I-language). A theory of universal grammar, like a particular proposed grammar, is true or false in whatever sense any scientific theory can be true or false. For our purposes we may accept the normal realist assumptions of the practicing scientist, in this connection. Whatever problems may arise are not specific to this enterprise, and are surely for better studied in connection with the more developed natural sciences.

Crucially, (I-)languages and S_0 are real entities, the basic objects of study for the science of language, though it may be possible to study more

complex abstractions, such as speech or language communities; any such further inquiry will surely have to presuppose grammars of (I-)language and universal grammar, and always has in practice, at least tacitly, even when this is explicitly denied, another confusion that I will not pursue here. An I-language — henceforth, simply “a language” — is the state attained by the language faculty, under certain external conditions. I doubt very much that it makes any sense to speak of a person as *learning* a language. Rather, a language grows in the mind/brain. Acquiring a language is less something that a child does than something that happens to the child, like growing arms rather than wings, or undergoing puberty at a certain stage of maturation. These processes take place in different ways depending on external events, but the basic lines of development are internally determined. The evidence seems to me overwhelming that this is true of language growth.

Let us now consider the question of knowledge. The language a person has acquired underlies a vast range of knowledge, both “knowledge-how” and “knowledge-that”. A person whose mind incorporates the language English (meaning, a particular I-language that falls within what is informally called “English”) knows how to speak and understand a variety of sentences, knows that certain sounds have certain meanings, and so on. These are typical cases of knowing-how and knowing-that, ordinary propositional knowledge in the latter case, and this of course does not exhaust the range of such knowledge. It seems entirely reasonable then to think of the language as a system that is internalized in the mind/brain, yielding specific cases of propositional knowledge and knowledge how to do so and so. We now have to consider at least three aspects of knowledge: (1) the internalized system of knowledge of the language, (2) knowing how to speak and understand, and (3) knowledge that sentences mean what they do (etc.).

It is common among philosophers, particularly those influenced by Wittgenstein, to hold that “knowledge of a language is an ability”, which can be exercised by speaking, understanding, reading, talking to oneself: “to know a language just is to have the ability to do these and similar things”², and indeed more generally, knowledge is a kind of ability. Some go further and hold that an ability is expressible in dispositional terms, so that language becomes, as Quine described it, “a complex of present dispositions to verbal behavior”. If we accept this further view, then two people who are disposed to say different things under given circumstances speak different languages, even if they are identical twins with exactly the same history, who speak the same language by any sensible criteria we might establish. There are so many well-known problems with this conception that I will simply drop it and consider the vaguer proposal that knowledge of language is a practical ability to speak and understand (Michael Dummett, Anthony Kenny, and others, in one or another form).

This radical departure from ordinary usage is, in my view, entirely unwarranted. To see how radical is the departure from ordinary usage, consider the consequences of accepting it, now using “ability” in the sense of ordinary usage. In the first place, ability can improve with no change of knowledge. Thus suppose someone takes a course in public speaking or in composition, improving his ability to speak and understand, but learning nothing new about his language. The language that the person speaks and understands is exactly what it was before, and the person’s knowledge of language has not changed, but his abilities have improved. Hence knowledge of language is not to be equated with the ability to speak, understand, etc.

Similarly, ability to use language can be impaired, and can even disappear, with no loss of knowledge of language at all. Suppose that a speaker of English suffers Parkinson’s disease, losing entirely the ability to speak, understand, etc. The person then does not have “the ability to do these and similar things”, and therefore does not have knowledge of English, as the term is defined by Kenny, Dummett and others. Suppose that use of the chemical L-Dopa can restore entirely the person’s ability, as has been claimed (it does not matter whether the facts just noted are accurate; since we are dealing with a conceptual question, it is enough that they could be, as is certainly the case). Now what has happened during the recovery of the ability? On the assumption in question, the person has recovered knowledge of English from scratch with a drug, after having totally lost that knowledge. Curiously, the person recovered knowledge of *English*, not of *Japanese*, though no evidence was available to choose between these outcomes; he regained knowledge of his original English with no experience at all. Had the person been a speaker of Japanese, he would have recovered Japanese with the same drug. Evidently, something remained fully intact while the ability was totally lost. In normal usage, as in our technical counterpart to it, we would say that what remained fully intact was “possession of the language”, knowledge of English, showing again that knowledge cannot be reduced to ability³.

Note that there are cases where we would say that a person retains an ability but is incapable of exercising it, say, a swimmer who cannot swim because his legs and arms are tied. But that is surely an entirely different kind of case than the one we are now considering, where the ability is lost, but the knowledge is retained.

To sustain the thesis that knowledge is ability, we would have to invent some new concept of ability, call it “P-ability” (philosopher’s ability). Then we could say that the person who improved his ability to speak without changing his knowledge retained this P-ability unchanged. And the person suffering from Parkinson’s disease retained his P-ability to speak (etc.) while entirely losing his ability to do so, in the normal sense of “ability”. Plainly this is pointless. The invented concept P-ability is invested with all the prop-

erties of knowledge, and diverges radically from the quite useful ordinary concept of ability, so nothing has been achieved except that we now mislead ourselves into believing that we have maintained the thesis that knowledge is ability. It is true that knowledge is P-ability, since we have defined “P-ability” to have the properties of knowledge, but that is hardly an interesting conclusion.

Exactly this tack is taken by Anthony Kenny, in the face of conceptual arguments such as those just reviewed. Thus in the case of the patient with Parkinson’s disease, Kenny says that he did indeed have the ability to use the language when he had no ability to use the language, thus shifting to “P-ability”, plainly, since the ability was totally lost⁴. Crucially, P-ability diverges radically from ability, and is like knowledge, as we can see from the fact that a person may have entirely lost the ability to speak and understand while entirely retaining the P-ability, can improve the ability with the P-ability unmodified, etc.

Kenny also assumes that there is a contradiction between my conclusion concerning the person who has lost the ability while retaining the knowledge and my statement elsewhere (which he accepts) that there might in principle be a “Spanish pill” that would confer knowledge of Spanish on a person who took it. There is no inconsistency. The issue in connection with aphasia or Parkinson’s disease has nothing to do with a pill for acquiring a particular language; rather, the point is that the person in the *Gedankenexperiment* reacquires ability to speak *exactly the language that he had* (knowledge of which he never lost); the same dose of L-dopa restores ability to speak English to the English speaker and ability to speak Japanese to the Japanese speaker; it is not an “English pill”. The same holds true of the person whose ability changes while his knowledge — or P-ability, if one prefers — remains constant.

It is curious that this attempt to maintain a clearly untenable thesis by inventing a new term “ability” that is used in the sense of “knowledge” and is radically different from “ability” in its normal sense is presented as in the spirit of Wittgenstein, who constantly inveighed against such procedures and argued that they are at the root of much philosophical error, as in the present case.

Note that essentially the same argument show that “knowing how” cannot be explained in terms of ability, unless we adopt the same pointless procedure just discussed. Suppose a person knows how to ride a bicycle, loses the ability to do so under some kind of brain injury, and then recovers it through administration of a drug, or when the effects of the injury recede. The person has made a transition from full ability, to no ability, to recovery of the original ability — not some other one. The argument is the same as before. Knowing how is not simply a matter of ability, nor, surely, is knowing that, contrary to much widely accepted doctrine. In fact, that is quite clear

from closer investigation of the concept “knowing how”. Rather, knowing how involves a crucial cognitive element, some internal representation of a system of knowledge⁵. Since this more general matter is not germane here, I will not pursue it.

Could we say, then, that knowing how to speak and understand a language is in no fundamental way different from knowing how to ride a bicycle, as is commonly alleged, so that we need not be driven to assume a mentally-represented system of knowledge in the case of language. There are at least two fundamental problems with this line of argument. First, knowing how in general involves a cognitive element, as just noted. Secondly, the “just like” argument is quite empty. We might as well say that there is no real problem in accounting for the ability that some people have to write brilliant poetry or wonderful quartets, or to discover deep theorems or scientific principles; it is just like knowing how to ride a bicycle. What possible point can there be to such proposals?

In any particular case, we have to discover what kind of cognitive structure underlies knowing how to do so-and-so or knowing that such-and-such⁶. In pursuing such inquiry, we rely entirely on “best theory” arguments, and we discover, not surprisingly, that very different kinds of systems, cognitive or other, are involved. To say that it is all just “knowing how”, hence unproblematic, is merely a form of anti-intellectualism, little more than an expression of lack of curiosity about features of the world, in this case, central features of human nature and human life.

In summary, to try to sustain the principle that knowing how to speak and understand a language reduces to a network of abilities, one has to use the term “ability” in some novel technical sense — in fact a sense invested with all the properties of knowledge. Plainly this is pointless.

A rather striking feature of the widespread conception of language as a system of abilities, or a habit system of some kind, or a complex of dispositions, is that it has been completely unproductive. It led precisely nowhere. One cannot point to a single result or discovery about language, even of the most trivial kind, that derives from this conception. Here one must be a bit more precise. There was, in fact, a discipline that professed this ideology, namely, American structural linguistics for many years. But the actual work carried out, and even the technical theories developed, departed from the ideology at every crucial point. Thus, there is no relation between, say, the procedures of phonemic analysis devised and the concept of language as a habit system⁷. This latter belief did influence applied disciplines such as language teaching, very much to their detriment. But linguistics itself was essentially unaffected, except insofar as it was impoverished in vision and concerns by the doctrine it professed.

One might draw an analogy to operationalism in the sciences. This ideology, widely professed at one time, undoubtedly had an influence in psychology. Namely, to the extent that it was followed in practice, it seriously impoverished the discipline. The principles were also professed in physics for a time, but I suspected that they had little impact there, since the people who professed the principles continued to do their work in utter contradiction to them, quite wisely.

A look at some concrete examples further illustrates the hopelessness of the attempt to explain knowledge in terms of ability. Consider the Spanish sentence (1):

(1) JUAN AFEITA A PEDRO.

We could replace *PEDRO* here by a reflexive, giving either (2*i*) or (2*ii*):

(2)(*i*) JUAN AFEITA A SÍ MISMO
(*ii*) JUAN AFEITA A SE.

In the latter case, a rule of cliticization applies, raising *SE* to the verb, giving (3), where *t* (“trace”) marks the position from which *SE* has moved and where it is interpreted as the direct object of *AFEITA*:

(3) JUAN SE AFEITA *t*.

Now consider the Spanish causative construction with a direct object or reflexive:

(4)(*i*) JUAN HIZO [AFEITAR A PEDRO]
(*ii*) JUAN HIZO [AFEITAR A SÍ MISMO]
(*iii*) JUAN HIZO [AFEITAR A SE].

In the last case, *SE* again moves, either to the verb *AFEITAR* or the verb *HIZO* (with perhaps some dialectal variation):

(5)(*i*) JUAN HIZO [AFEITARSE *t*]
(*ii*) JUAN SE HIZO [AFEITAR *t*].

Suppose now that we add the phrase *A LOS CHICOS* to (5*i*), yielding:

(6) JUAN HIZO [AFEITARSE *t* A LOS CHICOS]

Now the meaning changes: *SE* no longer refers to *JUAN* but to *LOS CHICOS*. Suppose we add *A LOS CHICOS* to (5*ii*), yielding:

(7) JUAN SE HIZO [AFEITAR t A LOS CHICOS].

This sentence is completely ungrammatical.

Similarly, suppose we precede the sentences (5) with the phrase *A QUIÉN*, forming a question, inducing a change of word order that we may put to the side:

(8)(i) A QUIÉN HIZO JUAN [AFEITARSE t]
 (ii) A QUIÉN SE HIZO JUAN [AFEITAR t]

In the first of these, *SE* once again does not refer to *JUAN*, as it does in (5*i*), but rather refers to *A QUIÉN*, or more accurately, to the variable *x* bound by the quasi-quantifier *A QUIÉN* in the mental representation:

(9) A QUIÉN HIZO JUAN [AFEITARSE t *x*]

Here the position of *x* is that of *LOS CHICOS* in (6). As for (8*ii*), it is completely ungrammatical, just like (7).

The question we now face is this: why does the meaning of “JUAN HIZO AFEITARSE” change completely when we add to it either “A LOS CHICOS” or “A QUIÉN”? And why does the sentence “JUAN SE HIZO AFEITAR” become completely ungrammatical when we add to it either of these phrases? If knowledge of language is a kind of ability, it must be that we do not have the ability to understand “JUAN HIZO AFEITARSE” with *SE* referring to *JUAN* when these phrases are added, and we do not have the ability to understand “JUAN SE HIZO AFEITAR” at all in these contexts. Perhaps we are too weak, or need to develop further skills. But this account is surely completely wide of the mark. The speaker of Spanish does not lack certain abilities. Rather, his or her system of mental computation simply works mechanically and virtually instantaneously to assign certain interpretations to certain representations in the mind/brain. The same is true of much simpler cases as well. If I know that the word *LIBRO* means “book” and not “table”, is my failure to have the word mean “table” a matter of lack of ability, because I am too weak or have not acquired some skill that I could try to perfect?

One can see, from such simple examples as these, why inquiry into language on the assumption that knowledge is a kind of ability has invariably proven so utterly fruitless, and why the assumption itself is a serious fallacy, which impedes the study of language and the study of knowledge to the extent that it is taken seriously.

Those examples suffice to illustrate some further points. The central problem of the theory of language is to explain how people can speak and understand new sentences, new in their experience or perhaps in the history of the language. The phenomenon is not an exotic one, but is the norm in the ordinary use of language, as Descartes and his followers stressed in their discussion of what we may call “the creative aspect of language use”, that is, the commonplace but often neglected fact that the normal use of language is unbounded in scope, free from identifiable stimulus control, coherent and appropriate to situation, evoking in listeners thoughts that they too might express. It is surprising how rarely the phenomenon was seriously addressed in the linguistics of the past century, until the mid-1950s at least, in part, perhaps, because of the conception of language as a system of habits, dispositions or abilities, Otto Jespersen being a rare and notable exception. When the question was addressed, the conventional answer was that new forms are produced and understood “by analogy” with familiar ones. The question then arises why the phrases “JUAN HIZO AFEITARSE” and “JUAN SE HIZO AFEITAR” are not understood “by analogy with themselves” when embedded in the context “— A LOS CHICOS” y “A QUIÉN —”. Surely, no analogy can be more straightforward than identity. The mind, however, chooses to ignore this straightforward analogy, and assigns a radically different interpretation to the identical forms. The situation is quite general, when we consider even the simplest cases, such as these.

The principle of computation that the mind is using in these cases can be discerned fairly easily. The reflexive form of *SE* searches for the nearest subject phrase in an appropriate abstract representation of the sentence, and associates itself with it. In the sentence “JUAN HIZO [AFEITARSE A LOS CHICOS]” the nearest subject is *LOS CHICOS*, so *SE* cannot be linked to *JUAN*. What about *JUAN SE HIZO [AFEITAR t A LOS CHICOS]*? By the principle just given *SE* links to *JUAN* but its trace *t* links to *A LOS CHICOS*. This yields a contradiction: the “chain” consisting of the pair (*SE*, *t*) is simply the abstract representation of the element *SE*, which receives its interpretation in the position of the trace; but this chain now is linked to two antecedents, *JUAN* and *LOS CHICOS*, and a reflexive, by virtue of its meaning, can have only one antecedent. So the sentence can receive no interpretation at all.

Note that this explanation relies on the assumption that the trace of *SE*, a so-called *empty category*, is really present in mental representation, “seen” by the mind just like any other element though not pronounced by the vocal tract. There is nothing problematic in this assumption, and it is justified by the fact that it yields an explanation for otherwise curious facts.

What about the sentence introduced by *A QUIÉN*? The same principle will account for the facts noted, as can easily be determined, if we assume

that the variable x is also present in the representation, appearing as the subject of the embedded sentence (10):

$$(10) \quad \text{AFEITAR}(SE) \ t x.$$

We therefore conclude that the mind “sees” two empty categories, the trace of *SE* and the variable bound by *A QUIÉN*, neither of which is produced by the vocal tract — though there is evidence that the variables can affect the pronunciation of neighboring elements. There are other examples that illustrate that the operator too can be an empty element.

Evidence of this sort very strongly supports the conclusion that there are empty categories with quite definite properties determined by the principles of universal grammar. The examples just cited cannot be translated into English, but English has other well-known phenomena, which cannot be translated into Spanish, which are explained by the same principles concerning reflexives and empty categories. The same is true of other languages, and there is reason to suppose that the principles are universal when precisely formulated, hence simply constituent elements of the initial state S_0 of the language faculty.

In the past few years it has been shown that quite a remarkable range of phenomena from typologically quite different languages can be explained on the assumption that the language faculty of the mind/brain carries out digital computations following very general principles, making use of representations of a precisely determined sort, including empty categories of several kinds. This work then provides evidence, quite strong evidence I believe, for some rather striking and surprising conclusions: that the language faculty, part of the mind/brain, is in crucial part a system of digital computation of a highly restricted character, with simple principles that interact to yield very intricate and complex results. This is a rather unexpected property of a biological system. One must be alert to the possibility that the conclusion is an artifact, resulting from our mode of analysis, but the evidence suggests quite strongly that the conclusion reflects reality.

As far as I am aware, there is only one other known biological phenomenon that shares the properties of discrete infinity exhibited by language, and that involves similar principles of digital computation: namely, the human number faculty, also apparently a species property, essentially common to the species and unique to it, and, like human language, unteachable to other organisms, which lack the requisite faculties. There are, for example, numerous animal communication systems, but they are invariably finite (the calls of apes) or continuous (the “language” of bees, continuous in whatever sense we can say this of a physical system; the human gestural system; etc.). Note that the difference between human languages and these communication

systems is not one of “more” or “less”, but one of difference in quality; indeed, it is doubtful that any sense can be given to the idea that human language *is* a communication system, though it can be used for communication, along with much else. These observations suggest that at some remote period of evolutionary history, the brain developed a certain capacity for digital computation, for employing recursive rules and associated mental representations, thus acquiring the basis for thought and language in the human sense, with the arithmetical capacity perhaps latent as a kind of abstraction from the language faculty, to be evoked when cultural conditions allowed, much later, in fact never in the case of some societies, so it appears. Notice that there is surely no reason to suppose that every trait is specifically selected.

The phenomena of the languages of the world appear to be highly diverse, but increasingly, it has been shown that over a large and impressive range they can be accounted for by the same principles, which yield highly varied results as the properties of lexical items vary from language to language. Thus in Spanish, there are clitic pronouns, including the reflexive, while in English there are not, so that the forms of English and Spanish, say in causative constructions, look quite different. But the principles that govern them appear to be essentially the same, their consequences differing by virtue of a lexical property of the pronominal system: in Spanish, but not English, there is a system of pronouns that are lexically marked as affixes, and therefore must attach to other elements. The manner in which these affixes attach, and the targets to which they adjoin, are determined by the very same principles that determine the formation of complex syntactic constructions such as operator-variable constructions and others, now applying to lexical items instead of to full phrases.

In other languages, many more terms are identified in the lexicon as affixes, and the same syntactic principles determine complex morphological forms that reflect in another way the same underlying and near universal underlying structures⁸. Thus in Japanese, the causative element is not a verb, like Spanish *HACER* or English *make*, but rather an affix, so a verb must move from the embedded clause to attach to it, yielding what appears to be a monoclausal causative as distinct from the English-Spanish biclausal causative; in Spanish too there is a reflection of the same process, when *SE* raises to the main verb in the sentence “JUAN SE HIZO AFEITAR”, as if *HIZO-AFEITAR* were a single word. The point is that as lexical items vary the very same principles determine a wide range of superficially different complex phenomena in typologically quite different languages.

The principles of universal grammar are fixed as constituent elements of the language faculty, but languages plainly differ. How do they differ? One way has already been noted: they differ in properties of lexical items, though here too the options are narrowly constrained by general principles.

Beyond that, it seems that the principles allow for a limited range of variation. That variation is limited has often been explicitly denied. The leading American linguist Edward Sapir held that languages can vary “without assignable limit”, and Martin Joos put forth what he called the “Boasian” view, referring to Franz Boas, one of the founders of modern linguistics: namely, that “languages could differ from each other without limit and in unpredictable ways”. Such views echo Nilliam Dwight Whitney, who greatly influenced Ferdinand de Saussure, and who emphasized “the infinite diversity of human speech”.

Such views perhaps appeared tenable in some form if one regarded language as a habit system, a network of practical abilities, a complex of dispositions, and the like. In that case, language would be constrained only by whatever general conditions constrain the development of abilities and habits in general, by what are sometimes called “generalized learning mechanisms”, if these exist. But this conception does not allow one even to approach the essential features of normal language use, as has been demonstrated beyond reasonable doubt in my view; and as already noted, the conception has been entirely unproductive.

Assuming without further discussion that this conception must be abandoned, the question of language variation will take on a new form, in the context of a general revision of the framework of inquiry into problems of natural language. A conceptual change of this nature was proposed about 30 years ago, reviving in a new form some long-forgotten approaches to the study of natural language. This rather sharp conceptual change underlies the research program that has been given the name “generative grammar”, referring to the fact that the grammar — or as we are now more properly calling it “the language” — generates an unbounded range of specific consequences, assigning a status to every expression and thus providing the mechanisms for the creative aspect of language use. The central questions of the study of language, conceived along the lines of the earlier discussion, now become the following:

- (11)(i) what is the system of knowledge attained by a person who speaks and understands a language?
- (ii) how is that knowledge acquired?
- (iii) how is that knowledge put to use?

The last question has two aspects, the production problem and the perception problem. The second question, how language is acquired, is a variant of what we might call “Plato’s problem”, raised for example when Socrates demonstrated that a slave boy with no training in geometry in fact knew geometry, perhaps the first psychological (thought-)experiment. The problem is not a

trivial one: people know a great deal more than can possibly be accounted for in terms of the standard paradigms of epistemology, language being a striking example. The production problem might be called “Descartes’ problem”, referring to one of the central Cartesian criteria for the existence of other minds: namely, when experiment demonstrates that another creature that resembles us exhibits the creative aspect of language use, then it would only be reasonable to attribute to the creature a mind like ours. In more recent years, a similar idea has been called “the Turing Test”. This problem, one aspect of more general problems concerning will and choice, remains beyond the scope of serious human inquiry in fact, and may be so in principle, rather as Descartes suggested. In any event, having nothing to say about it, I will put it aside, keeping just to the perception problem, or what is sometimes called “the parsing problem”.

These questions were posed as constituting the research program of generative grammar about 30 years ago, along with an argument to the effect that prevailing answers to them in terms of habit systems and the like were completely unacceptable for reasons already briefly discussed. What alternative, then, can we propose? I will keep to the terminology suggested above, departing from earlier usage.

The first proposal was that a language is a rule system, where the kinds of rules and their interrelations are specified by universal grammar. In one familiar conception, the rules included context-free rules, lexical rules, transformational rules, phonological rules (in a broad sense), and what were misleadingly called “rules of semantic interpretation” relating syntactic structures to representations in a system sometimes called “LF”, suggesting “logical form” but with certain qualifications. This term “rules of semantic interpretation” is misleading, as pointed out by David Lewis among others, because these rules relate syntactic objects, mental representations. They relate syntactic structures and LF-representations, which are syntactic objects. The term “semantics” should properly be restricted to the relation between language and the world, or to use of language, some might argue. The criticism is accurate, but it applies far more broadly. In fact, it applies in exactly the same form to what Lewis and others call “semantics”, where “meanings” are set-theoretic objects of some sort: models, “pictures”, situations and events, or whatever. These are mental representations⁹, not elements of the world, and the problem arises of how they are related to the world. It is often assumed that the relation is trivial, something like incorporation, so that it is unnecessary to provide a justification for these particular systems of mental representation, but it is easy to show that this cannot be true unless we trivialize our conception of what the world is by restricting attention to what Nelson Goodman calls “versions”, all mental representations, abandoning (perhaps as meaningless) the question of why one collection of “versions” is

jointly acceptable or “right” and others not, that is, not pursuing the common sense answer that certain versions are jointly “right” because of their accord with reality. But if we take this tack, which I do not suggest, semantics disappears and we are only studying various systems of mental representation. In fact, much of what is called “semantics” is really the study of the syntax of mental representations. It is a curious fact that those who correctly call their work in this area “syntax” are said to be avoiding semantics, while others who incorrectly describe their studies of syntax as “semantics” are said to be contributing to semantics¹⁰.

Adopting this conception of language, a language is a complex of rules of the permitted format, interconnected in a way permitted by universal grammar. In contrast to the conception of language in terms of habit systems or abilities, this was an extremely productive idea, which led quickly to a vast increase in the range of phenomena brought under investigation, with many discoveries about facts of language, even quite simple ones, that had never been noted or explored. Furthermore, the array of phenomena discovered and investigated were made intelligible at some level, by providing partial rule systems that accounted for their properties. The depth of explanation, however, could never really be very great. Even if appropriate rule systems could be constructed, and even if these systems were found to be restricted in type, we would always want to know why we have these kinds of rules and not others. Thus, languages typically have rules that allow the direct object of a verb to function as its subject, though it is still being interpreted as the object; but the converse property does not exist. Or consider again causative constructions, say, the form that we can give in abstract representation as (12) where the element CAUSE may be a word as in Spanish-English or on affix as in Japanese:

(12) problems CAUSE [that Y lies]

The principles of universal grammar permit a realization of this abstract form as something like (13), where CAUSE is an affix, or with CAUSE-*lie* associated in a closely linked verb sequence as in Spanish:

(13) problems CAUSE-*lie* Y.

But the form (14) does not underlie a possible realization as (15):

(14) [that Y lies CAUSE problems]

(15) Y CAUSE-*lie* problems.

Subject-object asymmetries of this sort are found very widely in language. They reflect in part the fact that subject-verb-object sentences are not treated in natural language as two-term relations as is familiar in logical analysis, but rather in the more traditional terms of Aristotelian logic and the universal grammar of the pre-modern period, as subject-predicate structures with a possibly complex predicate. In part, the asymmetries appear to follow from a newly discovered principle governing empty categories of the sort illustrated earlier. But whatever the explanation, problems of this nature abound, and an approach in terms of rule systems leaves them unsolved, except in a rather superficial way. From another point of view, there are simply too many possible rule systems, even when we constrain their form, and we thus do not achieve a convincing answer to our variant of Plato's problem.

Recognition of these facts has been at the core of the research program of the past 25 years. The natural approach has been to abandon rules in favor of general principles, so that the question of why we have one choice of rules rather than another simply does not arise. Thus if there are no rules for the formation of passive constructions, or interrogatives, or relative clauses, or phrase structure, and no rules that change grammatical functions such as causative and others then the question why we have certain rules, not others, does not arise. Increasingly, it has become clear that rules are simply epiphenomena, on a par with sentences in the sense that they are simply "projected" from the (I-)language, viewed in a certain way. But as distinct from sentences, which exist in mental representations and are realized in behavior, there is no reason to believe that rules of the familiar form exist at all; they have no status in linguistic theory and do not constitute part of mental representation or enter into mental computations, and we may safely abandon them, so it appears. We are left with general principles of universal grammar.

If there were only one possible human language, apart from lexical variety, we would then have a simple answer to our variant of Plato's problem: universal grammar permits only one realization apart from lexicon, and this is the language that people come to know when they acquire appropriate lexical items through experience in some manner. But clearly the variety of languages is greater than this, so this cannot be the complete story — though it is probably closer to true than has been thought in the past. Thus in languages such as English or Spanish, verbs and prepositions precede their objects, and the same is true of adjectives and nouns, as in such expressions as "ORGULLOSO DE SU HIJA" (where "SU HIJA" is the subject of "ORGULLOSO" with a semantically empty preposition *DE* introduced automatically as a kind of case-marker for reasons determined by universal grammar) and "TRADUCCIÓN DEL LIBRO", with a similar analysis. The categories noun, verb, adjective and preposition (more generally, apposition) are the *lexical categories*. The general principles of universal grammar determine the kinds of

phrases in which they appear as *heads*. The lexical entry itself determines the number and category of the *complements* of these heads and their semantic roles, and the general principles of phrase structure determine a limited range of other possibilities.

There is, however, an option left underdetermined by the principles of universal grammar. English and Spanish settle this option by placing the head invariably *before* its complements. We may say that they choose the “head-initial” value of the “head parameter”. In Japanese, in contrast, verbs, appositions, adjectives and nouns *follow* their complements. The range of phrase structures in the two languages is very similar, and accords with quite general principles of universal grammar, but the languages differ in one crucial choice of the head parameter: the language may choose either the “head-initial” or the “head-final” value of this parameter. In fact, this is only the simplest case, and there is a very limited range of further options depending on directionality of assignment of abstract case and semantic roles, a matter that has been explored by Hilda Koopman, Lisa Travis and others, but we may put these further complexities aside.

A crucial fact about the head parameter is that its value can be determined from very simple data. There is good reason to believe that this is true of all parameters; we must deal with the crucial and easily demonstrated fact that what a person knows is vastly underdetermined by available evidence, and that much of this knowledge is based on no direct evidence at all. Empty categories and their properties provide a dramatic example of this pervasive phenomenon, almost entirely ignored in earlier work. Thus a person is provided with no direct evidence about the position and various properties of elements that have no physical realization. There is little doubt that this problem of “poverty of stimulus” is in fact the norm rather than the exception. It must be, then, that the values of parameters are set by the kinds of simple data that are available to the child, and that the rich complex and highly articulated system of knowledge that arises, and is shared with others of somewhat different but equally impoverished experience, is determined in its basic features by the principles of the initial state S_0 of the language faculty. Languages may appear to differ, but they are cast to the same mold. One might draw an analogy to the biology of living organisms. Apparently, the biochemistry of life is quite similar from yeasts to humans, but small changes in timing of regulatory mechanisms of the cells and the like can yield what to us seem to be vast phenomenal differences, the difference between a whale and a butterfly, a human and a microbe, and so on. Viewed from an angel’s point of view, with numerous other possible though not actual sciences under consideration, all life might appear identical apart from trivialities. Similarly, from an angel’s point of view, all languages would appear identical, apart

from trivialities, their fundamental features determined by facts about human biology.

The language itself (again, as always, in the sense of I-language) may be regarded as nothing more than an array of choices for the various parameters, selected in accord with whatever options universal grammar permits. Since there are a finite number of parameters, each finite-valued (probably two-valued), it follows that there are a finite number of possible languages. One can see at once why questions concerning the formal properties of natural languages are largely irrelevant; there are few questions of mathematical interest to raise concerning finite sets.

Here a qualification is necessary. We are discussing what we might call “core language”, to be distinguished from a “periphery” of marked and specifically learned exceptions: irregular verbs, idioms, and the like. These may presumably vary without bound apart from time and memory limitations, though surely in a manner that is sharply constrained in type. It is the core language that is nothing other than an array of values for parameters. I assume, of course, that the distinction between core and periphery is a real world distinction, not a matter of convenience or pragmatic choice, except insofar as this is true of theories in chemistry and other branches of natural science, a consideration irrelevant here. For obvious reasons, the periphery is of much less interest for the basic psychological-biological questions to which linguistics is directed, if conceived along the lines of the previous discussion, and I will ignore it here.

Keeping to the core, then, there are finitely many possible languages. What a person knows, when that person speaks and understands a language, is a vocabulary and a particular array of values of parameters: an I-language. Once the parameters are set and lexical items acquired, the entire system functions, assigning a status to a vast range of expressions in a precise and explicit manner, even those that have never been heard or produced in the history of the language (and well beyond, as noted earlier). Others understand what we say, because they have the same biological nature and sufficiently similar experience with simple utterances.

Turning to Plato’s problem, a language is acquired by determining the values of the parameters of the initial state on the basis of simple data, and then the system of knowledge is represented in the mind/brain and is ready to function — though it might not function if the person lacks the ability to use it, perhaps because of some brain injury or the like. As for the parsing problem, it presumably should be solved along such lines as these: the hearer identifies words, and on the basis of their lexical properties, projects a syntactic structure as determined by principles of universal grammar and the values of the parameters. Connections and associations among these elements, including the empty categories that are forced to appear, are deter-

mined by other principles of universal grammar, perhaps parametrized. Thus given the sentence “A QUIÉN SE HIZO JUAN AFEITAR”, the mind of the speaker of Spanish automatically assigns a structure with two empty categories, one the subject of “AFEITAR”, another its object. Principles of universal grammar then produce a contradiction, in the manner informally described earlier, and the sentence receives no coherent interpretation, though of course it has a status: thus the Spanish speaker assigns to it a lexical and syntactic structure, and might even be able to “force” a certain meaning, if the sentence were produced by a foreigner, by me for example. A monolingual speaker of English will also assign a certain status to this expression, at least in some kind of phonetic representation, very likely considerably more.

The abandonment of rule systems in favor of principles-and-parameters approach, which has been gradually developing over the past 25 years and has been achieved to a substantial extent only in the past half-dozen years, has been extremely productive. It has, once again, led to a vast leap in empirical coverage, with entirely new empirical materials discovered in well-studied languages, and with languages of great typological variety incorporated within essentially the same framework. The depth of explanation has also advanced considerably, as it has become possible to explain, why there are processes described by certain rules but not others. The principles now being developed yield very sharp and surprising predictions about languages of varied types, predictions which sometimes prove accurate, and sometimes fail in highly instructive ways. My guess is that we are at the beginning of a radically new and highly productive phase in the study of language.

The principles-and-parameters approach yields a rather new way of thinking about questions of typology and comparative-historical linguistics. Consider again the analogy of speciation in biology, apparently small changes in the way fixed mechanisms function can produce large-scale phenomenal differences, yielding different species of organisms. In general, a slight change in the functioning of a rigidly structured and intricate system can yield very complex and surprising clusters of changes as its effects filter through the system. In the case of language, change of a single parameter may yield a cluster of differences which, on the surface, appear disconnected, as its effects filter through the invariant system of universal grammar. There is reason to believe that something of the sort is correct. Thus, among the Romance languages, French has a curious status. It differs from the other Romance languages in a cluster of properties, and it appears that these differences emerged fairly recently, and at about the same time. It may be that one parameter was changed — the null subject parameter that permits subject to be suppressed, some have speculated — yielding a cluster of other modifications through the mechanical working of the principles of universal grammar, and giving French something of the look of a Germanic language. At the

same time, French and Spanish share certain features distinguishing them from Italian, and there are numerous other complexities as we look at the actual languages, or “dialects” as they are called. Similarly, we find most remarkable similarities among languages that have no known historical connection, suggesting that they have simply set crucial parameters the same way. These are essentially new questions, which can now be seriously formulated for the first time and perhaps addressed.

As conceptions of language have changed over the years, so has the notion of what counts as a “real result”. Suppose we have some array of phenomena in some language. In the era of structural-descriptive linguistics, a result consisted in a useful arrangement of the data. As Zellig Harris put it in the major theoretical work of structural linguistics, a grammar provides a compact one-one representation of the phenomena in a corpus of data. Some, for example Roman Jakobson, went further in insisting on conformity to certain general laws, particularly in phonology, but in very limited ways.

Under the conception of language as a rule system, this would no longer count as a significant result; such a description poses rather than solves the problem at hand. Rather, it would be necessary to produce a rule system of the permitted format that predicts the data in question and in non-trivial cases, infinitely more. This is a much harder task, but not a hopeless one; there are many possible rule systems, and with effort, it is often possible to find one that satisfies the permitted format, if this is not too restricted.

Under the more recent principles-and-parameters approach, the task becomes harder still. A rule system is simply a description: it poses rather than solves the problem, and a “real result” consists of a demonstration that the phenomena under investigation, and countless others, can be explained by selecting properly the values of parameters in a fixed and invariant system of principles. This is a vastly harder problem, made still more difficult by the vast expansion of empirical materials in widely differing languages that have come to be partially understood, and to which any general theory must be responsible. Where the problem can be solved, we have results of real depth, well beyond anything imaginable earlier. It is an important fact that the problem is now intelligibly formulable, and that solutions are being produced over an interesting range, while efforts to pursue this inquiry are unearthing a large mass of new and unexplored phenomena in a wide variety of languages that pose new challenges, previously unknown.

This discussion has been based on the assumption that lexical items are somehow learned and available, suggesting that apart from parameter-setting, language acquisition as well as parsing and presumably the creative use of language (in the unlikely event that we can come to understand anything about this matter) are to a large extent determined by properties of the lexicon. But acquisition of lexical items poses Plato’s problem in a very sharp

form. As anyone who has tried to construct a dictionary or to work in descriptive semantics is aware, it is a very difficult matter to describe the meaning of a word, and such meanings have great intricacy and involve the most remarkable assumptions, even in the case of very simple concepts, such as what counts as a possible “thing”. At peak periods of language acquisition, children are “learning” many words a day, meaning that they are in effect learning words on a single exposure. This can only mean that the concepts are already available, with all or much of their intricacy and structure predetermined, and the child’s task is to assign labels to concepts, as might be done with very simple evidence.

Many have found this conclusion completely unacceptable, even absurd; it certainly departs radically from traditional views. Some, for example Hillary Putnam, have argued that it is entirely implausible to suppose that we have “an innate stock of notions” including *carburetor*, *bureaucrat*, etc.¹¹ If he were correct about this, it would not be particularly to the point, since the problem arises in a most serious way in connection with simple words such as “table”, “person”, “chase”, “persuade”, etc. But his argument for the examples he mentions is not compelling. It is that to have given us this innate stock of notions, “evolution would have had to be able to anticipate all the contingencies of *future* physical and cultural environments. Obviously it didn’t and couldn’t do this”. A very similar argument had long been accepted in immunology: namely, the number of antigens is so immense, including even artificially synthesized substances that had never existed in the world, that it was considered absurd to suppose that evolution had provided “an innate stock of antibodies”: rather, formation of antibodies must be a kind of “learning process” in which the antigens played an “instructive role”. But this assumption has been challenged, and is now widely assumed to be false. Niels Kaj Jerne won the Nobel Prize last year for his work challenging this idea, and upholding his own conception that an animal “cannot be stimulated to make specific antibodies, unless it has already made antibodies of this specificity before the antigen arrives”, so that antibody formation is a selective process in which the antigen plays a selective and amplifying role¹². Whether or not Jerne is correct, he certainly could be, and the same could be true in the case of word meanings, the argument being quite analogous.

Furthermore, there is good reason to suppose that the argument is at least in substantial measure correct, even for such words as *carburetor* and *bureaucrat*, which, in fact, pose the familiar problem of poverty of stimulus if we attend carefully to the enormous gap between what we know and the evidence on the basis of which we know it. The same is true of technical terms of science and mathematics, and it is quite surely the case for the terms of ordinary discourse. However surprising the conclusion may be that nature has

provided us with an innate stock of concepts, and that the child's task is to discover their labels, the empirical facts appear to leave open few other possibilities.

To the extent that anything is understood about lexical items and their nature, it seems that they are based on conceptual structures of a very specific and closely integrated type. It has been argued plausibly that concepts of a locational nature, including goal and source of action, object moved, place, etc., enter widely into lexical structure, often in quite abstract ways. In addition, notions like actor, recipient of action, event, intention, and others are pervasive elements of lexical structure, with their specific properties and permitted interrelations. Consider, say, the Spanish words *PERSEGUIR*, *PERSUADIR*. Like their English equivalents, these words clearly involve a reference to human intention. To chase someone is not only to follow him, but to follow him with the intent of staying on his path, perhaps to catch him. To persuade someone to do something is to cause him to decide or intend to do it; if he never decides or intends to do it, we have not succeeded in persuading him. Since these facts are known essentially without evidence, it must be that the child approaches language with an intuitive understanding of concepts involving intending, causation, goal of action, event, and so on, and places the words that are heard in a nexus that is permitted by the principles of universal grammar, which provide the framework for thought and language, and are common to human languages as conceptual systems of human life.

Notice further that we quite clearly have connections of meaning, analytic connections, in such cases as these; we have a rather clear distinction between truths of meaning and truths of fact. Thus, if John persuaded Bill to go to college, then Bill at some point decided or intended to go to college; otherwise, John did not persuade Bill to do so. This is a truth of meaning, not of fact. The *a priori* framework of human thought, within which language is acquired, provides necessary connections among concepts, reflected in connections of meaning among words, and more broadly, among expressions involving these words. Syntactic relations provide a rich array of further examples. It follows that one of the central conclusions of modern philosophy must be wrong: namely, the contention, often held to have been established by work of Quine and others, that one can make no principled distinction between questions of fact and questions of meaning, that it is a matter of more or less deeply held belief. Philosophers have, I think, been led to this erroneous conclusion, which is held to have undermined centuries of thought, by concentrating on an artificially narrow class of examples, in particular, on concepts that have little or no relational structure: such sentences as "cats are animals". Here, indeed, it is not easy to find evidence to decide whether the sentence is true as a matter of meaning or fact, and there has, again, been much inconclusive debate about the matter. When we turn to

more complex concepts with an inherent relational structure such as *PERSUADIR* or *PERSEGUIR*, or to more complex syntactic constructions such as “JUAN HIZO AFEITARSE A LOS CHICOS”, there seems little doubt that analytic connections are readily discerned.

Furthermore, the distinction, between truths of meaning and truths of fact now becomes an empirical matter, with a wide range of possible relevant evidence from the study of language acquisition, comparative linguistics, and other areas. There is, it seems clear, a rich conceptual structure determined by the initial state of the language faculty, waiting to be awakened by experience, much in accord with traditional rationalist conceptions and even the so-called “empiricist” thought of James Harris, David Hume and others.

I think we are forced to abandon many commonly accepted doctrines about language and knowledge. There is an innate structure that determines the framework within which thought and language develop, down to quite precise and intricate details. Language and thought are awakened in the mind, and follow a largely predetermined course, much like other biological properties. They develop in a way that provides a rich structure of truths of meaning. Our knowledge in these areas, and I believe elsewhere — even in science and mathematics — is not derived by induction, by applying reliable procedures, and so on; it is not grounded or based on “good reasons” in any useful sense of these notions. Rather, it grows in the mind, on the basis of our biological nature, triggered by appropriate experience, and in a limited way shaped by experience that settles options left open by the innate structure of mind. The result is an elaborate structure of cognitive systems, systems of knowledge and belief, that reflects the very nature of the human mind, a biological organ like others, with its scope and limits. This conclusion, which seems to me well-supported by the study of language and I suspect holds far more broadly, perhaps universally in domains of human thought, compels us to rethink fundamental assumptions of modern philosophy and of our general intellectual culture, including assumptions about scientific knowledge, mathematics, ethics, aesthetics, social theory and practice, and much else, questions too broad and far-reaching for me to try to address here, but questions that should, I think, be subjected to serious scrutiny from a point of view rather different from those that have conventionally been assumed.

*Department of Linguistics and Philosophy
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139, USA.*

NOTES

¹ For references, here and below, see my *Knowledge of Language: Its Nature, Origin and Use*, New York, Praeger, 1986.

² ANTHONY KENNY (1984), *The Legacy of Wittgenstein*, Oxford, Blackwell, p 138. Elsewhere Kenny speaks of “the futility of my attempt to separate knowledge of English from the ability to use — the mastery of — the language”. But to deny this identification of knowledge with ability is not to hold that knowledge can be “separated” from ability, whatever that means exactly.

³ Suppose that someone prefers to say that the knowledge of English was indeed lost, but that something else was retained. Then that “something else” is the only matter of interest for the new theory that will replace the old theory of knowledge, and the same conclusions follow: the only concept of significance, which plays the role of the now abandoned term “knowledge”, is this “possession of language” that cannot identified with ability. Clearly there is no point in these moves.

⁴ He also invests the invented concept of P-ability with curious properties, holding that had the patient not recovered with administration of the drug, he would not have had the P-ability when he lost the ability; but since the concept is invented, he may give it whatever properties he likes. To be precise, Kenny is not discussing the example given here but one that is identical in all relevant respects: an aphasic who loses all ability to use language and then recovers the ability in full when the effects of the injury recede. He also shifts from “ability” to “capacity” saying that when the person lacked the ability he had the capacity, thus using “capacity” in the sense of “knowledge” or “P-ability” but as distinct from “ability”, so that his claim about “knowledge” and “ability” collapses. In my *Rules and Representations*, Columbia, 1980, to which he refers in this connection, I pointed out that “capacity” is often used in a much looser sense than “ability”, so that a shift to “capacity” may disguise the falsehood of a characterization of knowledge in terms of ability. Kenny’s discussion is also marred by a fallacy of equivocation. Thus he notes that my usage of mentalistic terminology is quite different from his, but then criticizes my usage for “question-begging” because it would be nonsensical on his assumptions, which is correct but hardly relevant, since I was precisely challenging these assumptions, for the reasons reviewed here.

⁵ See my “Knowledge of Language”, in K. GUNDERSON, (ed.) (1975), *Language Mind and knowledge*, Minneapolis, U. of Minnesota Press.

⁶ If there is one. Note that I have not tried to establish that knowledge can never be reduced to ability, but rather that it cannot be in general, and in particular cannot be in the case of knowledge of language.

⁷ One cannot speak of strict inconsistency, since the concept of language as a habit system was regarded as a matter of fact, while the procedures of linguistic analysis as devised by many of the more sophisticated theorists were regarded as simply a device, one among many with no truth claim.

⁸ For very important recent discussion of this matter, see MARK BAKER (1985), *A Theory of Grammatical Function Changing*, MIT Ph.D. Dissertation.

⁹ At least, if we are doing C-linguistics, with empirical content. If not, then further clarification is required. The inquiry is in any event not semantics in the sense of empirical semantics, a study of relations between the language and something extralinguistic.

¹⁰ On a personal note, my work, from the beginning, has been largely concerned with the problem of developing linguistic theory so that the representations provided in particular languages will be appropriate for explaining how sentences are used and understood, but I have always called this “syntax”, as it is, even though the motivation is ultimately semantic; see, e.g., my *Logical Structure of Linguistic Theory*, 1955-6, published in part in 1975; New York, Plenum, U. of Chicago; *Syntactic Structures*, The Hague: Mouton, 1957. This work is correctly described as syntax, but it deals with questions that others incorrectly term “semantic”, and it is, I suspect, one crucial way to study semantics.

¹¹ See PUTNAM (1985), “Meaning and Our Mental Life”, manuscript.

¹² For discussion in a linguistic-cognitive context, see my *Rules and Representations*, New York, Columbia U. Press, 1980, 136f.; and JERNE’S Nobel Prize lecture (1985), “The Generative Grammar of the Immune System”, *Science* 229, pp.1057-9, Sept. 13.