



Noam Chomsky (1968)

Language and Mind

Linguistic Contributions to the Study of Mind (Future)

Source: *Language and Mind* publ. Harcourt Brace Jovanovich, Inc., 1968. One of the six lectures is reproduced here;

Transcribed: in 1998 by Andy Blunden, proofed and corrected February 2005.

In discussing the past, I referred to two major traditions that have enriched the study of language in their separate and very different ways; and in my last lecture, I tried to give some indication of the topics that seem on the immediate horizon today, as a kind of synthesis of philosophical grammar and structural linguistics begins to take shape. Each of the major traditions of study and speculation that I have been using as a point of reference was associated with a certain characteristic approach to the problems of mind; we might say, without distortion, that each evolved as a specific branch of the psychology of its time, to which it made a distinctive contribution.

It may seem a bit paradoxical to speak of structural linguistics in this way, given its militant anti-psychologism. But the paradox is lessened when we take note of the fact that this militant anti-psychologism is no less true of much of contemporary psychology itself, particularly of those branches that until a few years ago monopolised the study of use and acquisition of language. We live, after all, in the age of “behavioural science,” not of “the science of mind.” I do not want to read too much into a terminological innovation, but I think that there is some significance in

the ease and willingness with which modern thinking about man and society accepts the designation “behavioural science.” No sane person has ever doubted that behaviour provides much of the evidence for this study – all of the evidence, if we interpret “behaviour” in a sufficiently loose sense. But the term “behavioural science” suggests a not-so-subtle shift of emphasis toward the evidence itself and away from the deeper underlying principles and abstract mental structures that might be illuminated by the evidence of behaviour. It is as if natural science were to be designated “the science of meter readings.” What, in fact, would we expect of natural science in a culture that was satisfied to accept this designation for its activities?

Behavioural science has been much preoccupied with data and organisation of data, and it has even seen itself as a kind of technology of control of behaviour. Anti-mentalism in linguistics and in philosophy of language conforms to this shift of orientation. As I mentioned in my first lecture, I think that one major indirect contribution of modern structural linguistics results from its success in making explicit the assumptions of an anti-mentalistic, thoroughly operational and behaviourist approach to the phenomena of language. By extending this approach to its natural limits, it laid the groundwork for a fairly conclusive demonstration of the inadequacy of any such approach to the problems of mind.

More generally, I think that the long-range significance of the study of language lies in the fact that in this study it is possible to give a relatively sharp and clear formulation of some of the central questions of psychology and to bring a mass of evidence to bear on them. What is more, the study of language is, for the moment, unique in the combination it affords of richness of data and susceptibility to sharp formulation of basic issues.

It would, of course, be silly to try to predict the future of research, and it will be understood that I do not intend the subtitle of this lecture to be taken very seriously. Nevertheless, it is fair to suppose that the major contribution of the study of language will lie in the understanding it can provide as to the character of mental processes and the structures they form and manipulate. Therefore, instead of speculating on the likely course of research into the problems that are coming into focus today, I will concentrate here on some of the issues that arise when we try to develop the study of linguistic structure as a chapter of human psychology.

It is quite natural to expect that a concern for language will remain central to the study of human nature, as it has been in the past. Anyone concerned with the study of human nature and human capacities must somehow come to grips with the fact that all normal humans acquire language, whereas acquisition of even its barest rudiments is quite beyond the capacities of an otherwise intelligent ape – a fact that was emphasised, quite correctly, in Cartesian philosophy. It is widely thought that the extensive modern studies of animal communication challenge this classical view; and it is almost universally taken for granted that there exists a problem of explaining the “evolution” of human language from systems of animal communication. However, a careful look at recent studies of animal communication seems to me to provide little support for these assumptions. Rather, these studies simply bring out even more clearly the extent to which human language appears to be a unique phenomenon, without significant analogue in the animal world. If this is so, it is quite senseless to raise the problem of explaining the evolution of human language from more primitive systems of communication that appear at lower levels of intellectual capacity. The issue is important, and I would like to dwell on it for a moment.

The assumption that human language evolved from more primitive systems is developed in an interesting way by Karl Popper in his recently published Arthur Compton Lecture, "Clouds and Clocks." He tries to show how problems of freedom of will and Cartesian dualism can be solved by the analysis of this "evolution." I am not concerned now with the philosophical conclusions that he draws from this analysis, but with the basic assumption that there is an evolutionary development of language from simpler systems of the sort that one discovers in other organisms. Popper argues that the evolution of language passed through several stages, in particular a "lower stage" in which vocal gestures are used for expression of emotional state, for example, and a "higher stage" in which articulated sound is used for expression of thought – in Popper's terms, for description and critical argument. His discussion of stages of evolution of language suggests a kind of continuity, but in fact he establishes no relation between the lower and higher stages and does not suggest a mechanism whereby transition can take place from one stage to the next. In short, he gives no argument to show that the stages belong to a single evolutionary process. In fact, it is difficult to see what links these stages at all (except for the metaphorical use of the term "language"). There is no reason to suppose that the "gaps" are bridgeable. There is no more of a basis for assuming an evolutionary development of "higher" from "lower" stages, in this case, than there is for assuming an evolutionary development from breathing to walking; the stages have no significant analogy, it appears, and seem to involve entirely different processes and principles.

A more explicit discussion of the relation between human language and animal communication systems appears in a recent discussion by the comparative ethologist W. H. Thorpe. He points out that mammals other than man appear to lack the human ability to imitate sounds, and that one might therefore have expected birds (many of which have this ability to a

remarkable extent) to be “the group which ought to have been able to evolve language in the true sense, and not the mammals.” Thorpe does not suggest that human language “evolved” in any strict sense from simpler systems, but he does argue that the characteristic properties of human language can be found in animal communication systems, although “we cannot at the moment say definitely that they are all present in one particular animal.” The characteristics shared by human and animal language are the properties of being “purposive,” “syntactic,” and “propositional.” Language is purposive “in that there is nearly always in human speech a definite intention of getting something over to somebody else, altering his behaviour, his thoughts, or his general attitude toward a situation.” Human language is “Syntactic” in that an utterance is a performance with an internal organisation, with structure and coherence. It is “propositional” in that it transmits information. In this sense, then, both human language and animal communication are purposive, syntactic, and propositional.

All this may be true, but it establishes very little, since when we move to the level of abstraction at which human language and animal communication fall together, almost all other behaviour is included as well. Consider walking: Clearly, walking is purposive behaviour, in the most general sense of “purposive.” Walking is also “syntactic” in the sense just defined, as, in fact, Karl Lashley pointed out a long time ago in his important discussion of serial order in behaviour, to which I referred in the first lecture. Furthermore, it can certainly be informative; for example, I can signal my interest in reaching a certain goal by the speed or intensity with which I walk.

It is, incidentally, precisely in this manner that the examples of animal communication that Thorpe presents are “propositional.” He cites as an example the song of the European robin, in which the rate of alternation of high and low pitch signals the intention of the bird to defend its territory;

the higher the rate of alternation, the greater the intention to defend the territory. The example is interesting, but it seems to me to show very clearly the hopelessness of the attempt to relate human language to animal communication. Every animal communication system that is known (if we disregard some science fiction about dolphins) uses one of two basic principles: Either it consists of a fixed, finite number of signals, each associated with a specific range of behaviour or emotional state, as is illustrated in the extensive primate studies that have been carried out by Japanese scientists for the past several years; or it makes use of a fixed, finite number of linguistic dimensions, each of which is associated with a particular nonlinguistic dimension in such a way that selection of a point along the linguistic dimension determines and signals a certain point along the associated nonlinguistic dimension. The latter is the principle realised in Thorpe's bird-song example. Rate of alternation of high and low pitch is a linguistic dimension correlated with the nonlinguistic dimension of intention to defend a territory. The bird signals its intention to defend a territory by selecting a correlated point along the linguistic dimension of pitch alternation – I use the word “select” loosely, of course. The linguistic dimension is abstract, but the principle is clear. A communication system of the second type has an indefinitely large range of potential signals, as does human language. The mechanism and principle, however, are entirely different from those employed by human language to express indefinitely many new thoughts, intentions, feelings, and so on. It is not correct to speak of a “deficiency” of the animal system, in terms of range of potential signals; rather the opposite, since the animal system admits in principle of continuous variation along the linguistic dimension (insofar as it makes sense to speak of “continuity” in such a case), whereas human language is discrete. Hence, the issue is not one of “more” or “less,” but rather of an entirely different principle of organisation. When I make some arbitrary statement in a human language – say, that “the rise of supranational corporations poses new dangers for

human freedom” – I am not selecting a point along some linguistic dimension that signals a corresponding point along an associated nonlinguistic dimension, nor am I selecting a signal from a finite behavioural repertoire, innate or learned.

Furthermore, it is wrong to think of human use of language as characteristically informative, in fact or in intention. Human language can be used to inform or mislead, to clarify one’s own thoughts or to display one’s cleverness, or simply for play. If I speak with no concern for modifying your behaviour or thoughts, I am not using language any less than if I say exactly the same things with such intention. If we hope to understand human language and the psychological capacities on which it rests, we must first ask what it is, not how or for what purposes it is used. When we ask what human language is, we find no striking similarity to animal communication systems. There is nothing useful to be said about behaviour or thought at the level of abstraction at which animal and human communication fall together. The examples of animal communication that have been examined to date do share many of the properties of human gestural systems, and it might be reasonable to explore the possibility of direct connection in this case. But human language, it appears, is based on entirely different principles. This, I think, is an important point, often overlooked by those who approach human language as a natural, biological phenomenon; in particular, it seems rather pointless, for these reasons, to speculate about the evolution of human language from simpler systems – perhaps as absurd as it would be to speculate about the “evolution” of atoms from clouds of elementary particles.

As far as we know, possession of human language is associated with a specific type of mental organisation, not simply a higher degree of intelligence. There seems to be no substance to the view that human language is simply a more complex instance of something to be found

elsewhere in the animal world. This poses a problem for the biologist, since, if true, it is an example of true “emergence” – the appearance of a qualitatively different phenomenon at a specific stage of complexity of organisation. Recognition of this fact, though formulated in entirely different terms, is what motivated much of the classical study of language by those whose primary concern was the nature of mind. And it seems to me that today there is no better or more promising way to explore the essential and distinctive properties of human intelligence than through the detailed investigation of the structure of this unique human possession. A reasonable guess, then, is that if empirically adequate generative grammars can be constructed and the universal principles that govern their structure and organisation determined, then this will be an important contribution to human psychology, in ways to which I will turn directly, in detail.

In the course of these lectures I have mentioned some of the classical ideas regarding language structure and contemporary efforts to deepen and extend them. It seems clear that we must regard linguistic competence – knowledge of a language – as an abstract system underlying behaviour, a system constituted by rules that interact to determine the form and intrinsic meaning of a potentially infinite number of sentences. Such a system – a generative grammar – provides an explication of the Humboldtian idea of “form of language,” which in an obscure but suggestive remark in his great posthumous work, *Über die Verschiedenheit des Menschlichen Sprachbaues*, Humboldt defines as “that constant and unvarying system of processes underlying the mental act of raising articulated structurally organised signals to an expression of thought.” Such a grammar defines a language in the Humboldtian sense, namely as “a recursively generated system, where the laws of generation are fixed and invariant, but the scope and the specific manner in which they are applied remain entirely unspecified.”

In each such grammar there are particular, idiosyncratic elements, selection of which determines one specific human language; and there are general universal elements, conditions on the form and organisation of any human language, that form the subject matter for the study of “universal grammar.” Among the principles of universal grammar are those I discussed in the preceding lecture – for example, the principles that distinguish deep and surface structure and that constrain the class of transformational operations that relate them. Notice, incidentally, that the existence of definite principles of universal grammar makes possible the rise of the new field of mathematical linguistics, a field that submits to abstract study the class of generative systems meeting the conditions set forth in universal grammar. This inquiry aims to elaborate the formal properties of any possible human language. The field is in its infancy; it is only in the last decade that the possibility of such an enterprise has been envisioned. It has some promising initial results, and it suggests one possible direction for future research that might prove to be of great importance. Thus, mathematical linguistics seems for the moment to be in a uniquely favourable position, among mathematical approaches in the social and psychological sciences, to develop not simply as a theory of data, but as the study of highly abstract principles and structures that determine the character of human mental processes. In this case, the mental processes in question are those involved in the organisation of one specific domain of human knowledge, namely knowledge of language.

The theory of generative grammar, both particular and universal, points to a conceptual lacuna in psychological theory that I believe is worth mentioning. Psychology conceived as “behavioural science” has been concerned with behaviour and acquisition or control of behaviour. It has no concept corresponding to “competence,” in the sense in which competence is characterised by a generative grammar. The theory of learning has limited itself to a narrow and surely inadequate concept of

what is learned – namely a system of stimulus-response connections, a network of associations, a repertoire of behavioural items, a habit hierarchy, or a system of dispositions to respond in a particular way under specifiable stimulus conditions. Insofar as behavioural psychology has been applied to education or therapy, it has correspondingly limited itself to this concept of “what is learned.” But a generative grammar cannot be characterised in these terms. What is necessary, in addition to the concept of behaviour and learning, is a concept of what is learned – a notion of competence – that lies beyond the conceptual limits of behaviourist psychological theory. Like much of modern linguistics and modern philosophy of language, behaviourist psychology has quite consciously accepted methodological restrictions that do not permit the study of systems of the necessary complexity and abstractness. One important future contribution of the study of language to general psychology may be to focus attention on this conceptual gap and to demonstrate how it may be filled by the elaboration of a system of underlying competence in one domain of human intelligence.

There is an obvious sense in which any aspect of psychology is based ultimately on the observation of behaviour. But it is not at all obvious that the study of learning should proceed directly to the investigation of factors that control behaviour or of conditions under which a “behavioural repertoire” is established. It is first necessary to determine the significant characteristics of this behavioural repertoire, the principles on which it is organised. A meaningful study of learning can proceed only after this preliminary task has been carried out and has led to a reasonably well-confirmed theory of underlying competence – in the case of language, to the formulation of the generative grammar that underlies the observed use of language. Such a study will concern itself with the relation between the data available to the organism and the competence that it acquires; only to the extent that the abstraction to competence has been successful – in the

case of language, to the extent that the postulated grammar is “descriptively adequate” in the sense described in Lecture 2 – can the investigation of learning hope to achieve meaningful results. If, in some domain, the organisation of the behavioural repertoire is quite trivial and elementary, then there will be little harm in avoiding the intermediate stage of theory construction, in which we attempt to characterise accurately the competence that is acquired. But one cannot count on this being the case, and in the study of language it surely is not the case. With a richer and more adequate characterisation of “what is learned” – of the underlying competence that constitutes the “final state” of the organism being studied – it may be possible to approach the task of constructing a theory of learning that will be much less restricted in scope than modern behavioural psychology has proved to be. Surely it is pointless to accept methodological strictures that preclude such an approach to problems of learning.

Are there other areas of human competence where one might hope to develop a fruitful theory, analogous to generative grammar? Although this is a very important question, there is very little that can be said about it today. One might, for example, consider the problem of how a person comes to acquire a certain concept of three-dimensional space, or an implicit “theory of human action,” in similar terms. Such a study would begin with the attempt to characterise the implicit theory that underlies actual performance and would then turn to the question of how this theory develops under the given conditions of time and access to data, that is, in what way the resulting system of beliefs is determined by the interplay of available data, “heuristic procedures,” and the innate schematism that restricts and conditions the form of the acquired system. At the moment, this is nothing more than a sketch of a program of research.

There have been some attempts to study the structure of other, language-like systems – the study of kinship systems and folk taxonomies

comes to mind, for example. But so far, at least, nothing has been discovered that is even roughly comparable to language in these domains. No one, to my knowledge, has devoted more thought to this problem than Lévi-Strauss. For example, his recent book on the categories of primitive mentality is a serious and thoughtful attempt to come to grips with this problem. Nevertheless, I do not see what conclusions can be reached from a study of his materials beyond the fact that the savage mind attempts to impose some organisation on the physical world – that humans classify, if they perform any mental acts at all. Specifically, Lévi-Strauss's well-known critique of totemism seems to reduce to little more than this conclusion.

Lévi-Strauss models his investigations quite consciously on structural linguistics, particularly on the work of Troubetzkoy and Jakobson. He repeatedly and quite correctly emphasises that one cannot simply apply procedures analogous to those of phonemic analysis to subsystems of society and culture. Rather, he is concerned with structures “where they may be found ... in the kinship system, political ideology, mythology, ritual, art,” and so on, and he wishes to examine the formal properties of these structures in their own terms. But several reservations are necessary when structural linguistics is used as a model in this way. For one thing, the structure of a phonological system is of very little interest as a formal object; there is nothing of significance to be said, from a formal point of view, about a set of forty-odd elements cross-classified in terms of eight or ten features. The significance of structuralist phonology, as developed by Troubetzkoy, Jakobson, and others, lies not in the formal properties of phonemic systems but in the fact that a fairly small number of features that can be specified in absolute, language-independent terms appear to provide the basis for the organisation of all phonological systems. The achievement of structuralist phonology was to show that the phonological rules of a great variety of languages apply to classes of elements that can

be simply characterised in terms of these features; that historical change affects such classes in a uniform way; and that the organisation of features plays a basic role in the use and acquisition of language. This was a discovery of the greatest importance, and it provides the groundwork for much of contemporary linguistics. But if we abstract away from the specific universal set of features and the rule systems in which they function, little of any significance remains.

Furthermore, to a greater and greater extent, current work in phonology is demonstrating that the real richness of phonological systems lies not in the structural patterns of phonemes but rather in the intricate systems of rules by which these patterns are formed, modified, and elaborated. The structural patterns that arise at various stages of derivation are a kind of epiphenomenon. The system of phonological rules makes use of the universal features in a fundamental way, but it is the properties of the systems of rules, it seems to me, that really shed light on the specific nature of the organisation of language. For example, there appear to be very general conditions, such as the principle of cyclic ordering (discussed in the preceding lecture) and others that are still more abstract, that govern the application of these rules, and there are many interesting and unsolved questions as to how the choice of rules is determined by intrinsic, universal relations among features. Furthermore, the idea of a mathematical investigation of language structures, to which Lévi-Strauss occasionally alludes, becomes meaningful only when one considers systems of rules with infinite generative capacity. There is nothing to be said about the abstract structure of the various patterns that appear at various stages of derivation. If this is correct, then one cannot expect structuralist phonology, in itself, to provide a useful model for investigation of other cultural and social systems.

In general, the problem of extending concepts of linguistic structure to other cognitive systems seems to me, for the moment, in not too promising a state, although it is no doubt too early for pessimism.

Before turning to the general implications of the study of linguistic competence and, more specifically, to the conclusions of universal grammar, it is well to make sure of the status of these conclusions in the light of current knowledge of the possible diversity of language. In my first lecture, I quoted the remarks of William Dwight Whitney about what he referred to as “the infinite diversity of human speech,” the boundless variety that, he maintained, undermines the claims of philosophical grammar to psychological relevance.

Philosophical grammarians had typically maintained that languages vary little in their deep structures, though there may be wide variability in surface manifestations. Thus there is, in this view, an underlying structure of grammatical relations and categories, and certain aspects of human thought and mentality are essentially invariant across languages, although languages may differ as to whether they express the grammatical relations formally by inflection or word order, for example. Furthermore, an investigation of their work indicates that the underlying recursive principles that generate deep structure were assumed to be restricted in certain ways – for example, by the condition that new structures are formed only by the insertion of new “propositional content,” new structures that themselves correspond to actual simple sentences, in fixed positions in already formed structures. Similarly, the grammatical transformations that form surface structures through reordering, ellipsis, and other formal operations must themselves meet certain fixed general conditions, such as those discussed in the preceding lecture. In short, the theories of philosophical grammar, and the more recent elaborations of these theories, make the assumption that languages will differ very little, despite considerable diversity in superficial realisation, when we discover

their deeper structures and unearth their fundamental mechanisms and principles.

It is interesting to observe that this assumption persisted even through the period of German romanticism, which was, of course, much preoccupied with the diversity of cultures and with the many rich possibilities for human intellectual development. Thus, Wilhelm von Humboldt, who is now best remembered for his ideas concerning the variety of languages and the association of diverse language structures with divergent “world-views,” nevertheless held firmly that underlying any human language we will find a system that is universal, that simply expresses man’s unique intellectual attributes. For this reason, it was possible for him to maintain the rationalist view that language is not really learned – certainly not taught – but rather develops “from within,” in an essentially predetermined way, when the appropriate environmental conditions exist. One cannot really teach a first language, he argued, but can only “provide the thread along which it will develop of its own accord,” by processes more like maturation than learning. This Platonistic element in Humboldt’s thought is a pervasive one; for Humboldt, it was as natural to propose an essentially Platonistic theory of “learning” as it was for Rousseau to found his critique of repressive social institutions on a conception of human freedom that derives from strictly Cartesian assumptions regarding the limitations of mechanical explanation. And in general it seems appropriate to construe both the psychology and the linguistics of the romantic period as in large part a natural outgrowth of rationalist conceptions.

The issue raised by Whitney against Humboldt and philosophical grammar in general is of great significance with respect to the implications of linguistics for general human psychology. Evidently, these implications can be truly far-reaching only if the rationalist view is essentially correct, in which case the structure of language can truly serve

as a “mirror of mind,” in both its particular and its universal aspects. It is widely believed that modern anthropology has established the falsity of the assumptions of the rationalist universal grammarians by demonstrating through empirical study that languages may, in fact, exhibit the widest diversity. Whitney’s claims regarding the diversity of languages are reiterated throughout the modern period; Martin Joos, for example, is simply expressing the conventional wisdom when he takes the basic conclusion of modern anthropological linguistics to be that “languages can differ without limit as to either extent or direction.”

The belief that anthropological linguistics has demolished the assumptions of universal grammar seems to me to be quite false in two important respects. First, it misinterprets the views of classical rationalist grammar, which held that languages are similar only at the deeper level, the level at which grammatical relations are expressed and at which the processes that provide for the creative aspect of language use are to be found. Second, this belief seriously misinterprets the findings of anthropological linguistics, which has, in fact, restricted itself almost completely to fairly superficial aspects of language structure.

To say this is not to criticise anthropological linguistics, a field that is faced with compelling problems of its own – in particular, the problem of obtaining at least some record of the rapidly vanishing languages of the primitive world. Nevertheless, it is important to bear in mind this fundamental limitation on its achievements in considering the light it can shed on the theses of universal grammar. Anthropological studies (like structural linguistic studies in general) do not attempt to reveal the underlying core of generative processes in language – that is, the processes that determine the deeper levels of structure and that constitute the systematic means for creating ever novel sentence types. Therefore, they obviously cannot have any real bearing on the classical assumption that these underlying generative processes vary only slightly from

language to language. In fact, what evidence is now available suggests that if universal grammar has serious defects, as indeed it does from a modern point of view, then these defects lie in the failure to recognise the abstract nature of linguistic structure and to impose sufficiently strong and restrictive conditions on the form of any human language. And a characteristic feature of current work in linguistics is its concern for linguistic universals of a sort that can only be detected through a detailed investigation of particular languages, universals governing properties of language that are simply not accessible to investigation within the restricted framework that has been adopted, often for very good reasons, within anthropological linguistics.

I think that if we contemplate the classical problem of psychology, that of accounting for human knowledge, we cannot avoid being struck by the enormous disparity between knowledge and experience – in the case of language, between the generative grammar that expresses the linguistic competence of the native speaker and the meagre and degenerate data on the basis of which he has constructed this grammar for himself. In principle the theory of learning should deal with this problem; but in fact it bypasses the problem, because of the conceptual gap that I mentioned earlier. The problem cannot even be formulated in any sensible way until we develop the concept of competence, alongside the concepts of learning and behaviour, and apply this concept in some domain. The fact is that this concept has so far been extensively developed and applied only in the study of human language. It is only in this domain that we have at least the first steps toward an account of competence, namely the fragmentary generative grammars that have been constructed for particular languages. As the study of language progresses, we can expect with some confidence that these grammars will be extended in scope and depth, although it will hardly come as a surprise if the first proposals are found to be mistaken in fundamental ways.

Insofar as we have a tentative first approximation to a generative grammar for some language, we can for the first time formulate in a useful way the problem of origin of knowledge. In other words, we can ask the question, What initial structure must be attributed to the mind that enables it to construct such a grammar from the data of sense? Some of the empirical conditions that must be met by any such assumption about innate structure are moderately clear. Thus, it appears to be a species-specific capacity that is essentially independent of intelligence, and we can make a fairly good estimate of the amount of data that is necessary for the task to be successfully accomplished. We know that the grammars that are in fact constructed vary only slightly among speakers of the same language, despite wide variations not only in intelligence but also in the conditions under which language is acquired. As participants in a certain culture, we are naturally aware of the great differences in ability to use language, in knowledge of vocabulary, and so on that result from differences in native ability and from differences in conditions of acquisition; we naturally pay much less attention to the similarities and to common knowledge, which we take for granted. But if we manage to establish the requisite psychic distance, if we actually compare the generative grammars that must be postulated for different speakers of the same language, we find that the similarities that we take for granted are quite marked and that the divergences are few and marginal. What is more, it seems that dialects that are superficially quite remote, even barely intelligible on first contact, share a vast central core of common rules and processes and differ very slightly in underlying structures, which seem to remain invariant through long historical eras. Furthermore, we discover a substantial system of principles that do not vary among languages that are, as far as we know, entirely unrelated.

The central problems in this domain are empirical ones that are, in principle at least, quite straightforward, difficult as they may be to solve in

a satisfactory way. We must postulate an innate structure that is rich enough to account for the disparity between experience and knowledge, one that can account for the construction of the empirically justified generative grammars within the given limitations of time and access to data. At the same time, this postulated innate mental structure must not be so rich and restrictive as to exclude certain known languages. There is, in other words, an upper bound and a lower bound on the degree and exact character of the complexity that can be postulated as innate mental structure. The factual situation is obscure enough to leave room for much difference of opinion over the true nature of this innate mental structure that makes acquisition of language possible. However, there seems to me to be no doubt that this is an empirical issue, one that can be resolved by proceeding along the lines that I have just roughly outlined.

My own estimate of the situation is that the real problem for tomorrow is that of discovering an assumption regarding innate structure that is sufficiently rich, not that of finding one that is simple or elementary enough to be “plausible.” There is, as far as I can see, no reasonable notion of “plausibility,” no *a priori* insight into what innate structures are permissible, that can guide the search for a “sufficiently elementary assumption.” It would be mere dogmatism to maintain without argument or evidence that the mind is simpler in its innate structure than other biological systems, just as it would be mere dogmatism to insist that the mind’s organisation must necessarily follow certain set principles, determined in advance of investigation and maintained in defiance of any empirical findings. I think that the study of problems of mind has been very definitely hampered by a kind of apriorism with which these problems are generally approached. In particular, the empiricist assumptions that have dominated the study of acquisition of knowledge for many years seem to me to have been adopted quite without warrant

and to have no special status among the many possibilities that one might imagine as to how the mind functions.

In this connection, it is illuminating to follow the debate that has arisen since the views I have just sketched were advanced a few years ago as a program of research – I should say, since this position was resurrected, because to a significant extent it is the traditional rationalist approach, now amplified and sharpened and made far more explicit in terms of the tentative conclusions that have been reached in the recent study of linguistic competence. Two outstanding American philosophers, Nelson Goodman and Hilary Putnam, have made recent contributions to this discussion – both misconceived, in my opinion, but instructive in the misconceptions that they reveal.

Goodman's treatment of the question suffers first from an historical misunderstanding and second from a failure to formulate correctly the exact nature of the problem of acquisition of knowledge. His historical misunderstanding has to do with the issue between Locke and whomever Locke thought he was criticising in his discussion of innate ideas. According to Goodman, "Locke made ... acutely clear" that the doctrine of innate ideas is "false or meaningless." In fact, however, Locke's critique had little relevance to any familiar doctrine of the seventeenth century. The arguments that Locke gave were considered and dealt with in quite a satisfactory way in the earliest seventeenth-century discussions of innate ideas, for example those of Lord Herbert and Descartes, both of whom took for granted that the system of innate ideas and principles would not function unless appropriate stimulation took place. For this reason, Locke's arguments, none of which took cognisance of this condition, are without force; for some reason, he avoided the issues that had been discussed in the preceding half-century. Furthermore, as Leibnitz observed, Locke's willingness to make use of a principle of "reflection" makes it almost impossible to distinguish his approach from

that of the rationalists, except for his failure to take even those steps suggested by his predecessors toward specifying the character of this principle.

But, historical issues aside, I think that Goodman misconstrues the substantive problem as well. He argues that first-language learning poses no real problem, because prior to first-language learning the child has already acquired the rudiments of a symbolic system in his ordinary dealings with the environment. Hence, first-language learning is analogous to second-language learning in that the fundamental step has already been taken, and details can be elaborated within an already existing framework. This argument might have some force if it were possible to show that the specific properties of grammar – say, the distinction of deep and surface structure, the specific properties of grammatical transformations, the principles of rule ordering, and so on – were present in some form in these already acquired prelinguistic “symbolic systems.” But since there is not the slightest reason to believe that this is so, the argument collapses. It is based on an equivocation similar to that discussed earlier in connection with the argument that language evolved from animal communication. In that case, as we observed, the argument turned on a metaphorical use of the term “language.” In Goodman’s case, the argument is based entirely on a vague use of the term “symbolic system,” and it collapses as soon as we attempt to give this term a precise meaning. If it were possible to show that these prelinguistic symbolic systems share certain significant properties with natural language, we could then argue that these properties of natural language are acquired by analogy. Of course, we would then face the problem of explaining how the prelinguistic symbolic systems developed these properties. But since no one has succeeded in showing that the fundamental properties of natural language – those discussed in Lecture 2,

for example – appear in prelinguistic symbolic systems or any others, the latter problem does not arise.

According to Goodman, the reason why the problem of second-language learning is different from that of first-language learning is that “once one language is available,” it “can be used for giving explanation and instruction.” He then goes on to argue that “acquisition of an initial language is acquisition of a secondary symbolic system” and is quite on a par with normal second-language acquisition. The primary symbolic systems to which he refers are “rudimentary-prelinguistic symbolic systems in which gestures and sensory and perceptual occurrences of all sorts function as signs.” But evidently these prelinguistic symbolic systems cannot be “used for giving explanation and instruction” in the way a first language can be used in second-language instruction. Therefore, even on his own grounds, Goodman’s argument is incoherent.

Goodman maintains that “the claim we are discussing cannot be experimentally tested even when we have an acknowledged example of a ‘bad’ language” and that “the claim has not even been formulated to the extent of citation of a single general property of ‘bad’ languages.” The first of these conclusions is correct, in his sense of “experimental test,” namely a test in which we “take an infant at birth, isolate it from all the influences of our language-bound culture, and attempt to inculcate it with one of the ‘bad’ artificial languages.” Obviously this is not feasible. But there is no reason why we should be dismayed by the impossibility of carrying out such a test as this. There are many other ways, for example, those discussed in Lecture 2 and the references cited there – in which evidence can be obtained concerning the properties of grammars and conclusions regarding the general properties of such grammars can be put to empirical test. Any such conclusion immediately specifies, correctly or incorrectly, certain properties of “bad” languages. Since there are dozens of papers and books that attempt to formulate such properties, his second

claim, that not “a single general property of ‘bad’ languages” has been formulated, is rather surprising. One might try to show that these attempts are misguided or questionable, but one can hardly maintain seriously that they do not exist. Any formulation of a principle of universal grammar makes a strong empirical claim, which can be falsified by finding counterinstances in some human language, along the lines of the discussion in Lecture 2. In linguistics, as in any other field, it is only in such indirect ways as this that one can hope to find evidence bearing on non-trivial hypotheses. Direct experimental tests of the sort that Goodman mentions are rarely possible, a matter that may be unfortunate but is nevertheless characteristic of most research.

At one point Goodman remarks, correctly, that even though “for certain remarkable facts I have no alternative explanation ... that alone does not dictate acceptance of whatever theory may be offered; for the theory might be worse than none. Inability to explain a fact does not condemn me to accept an intrinsically repugnant and incomprehensible theory.” But now consider the theory of innate ideas that Goodman regards as “intrinsically repugnant and incomprehensible.” Notice, first, that the theory is obviously not “incomprehensible,” on his terms. Thus he appears to be willing, in this article, to accept the view that in some sense the mature mind contains ideas; it is obviously not “incomprehensible,” then, that some of these ideas are “implanted in the mind as original equipment,” to use his phraseology. And if we turn to the actual doctrine as developed in rationalist philosophy, rather than Locke’s caricature, the theory becomes even more obviously comprehensible. There is nothing incomprehensible in the view that stimulation provides the occasion for the mind to apply certain innate interpretive principles, certain concepts that proceed from “the power of understanding” itself, from the faculty of thinking rather than from external objects directly. To take an example from Descartes (*Reply to Objections*, V):

When first in infancy we see a triangular figure depicted on paper, this figure cannot show us how a real triangle ought to be conceived in the way in which geometers consider it, because the true triangle is contained in this figure, just as the statue of Mercury is contained in a rough block of wood. But because we already possess within us the idea of a true triangle, and it can be more easily conceived by our mind than the more complex figure of the triangle drawn on paper, we, therefore, when we see the composite figure, apprehend not it itself, but rather the authentic triangle.

In this sense the idea of a triangle is innate. Surely the notion is comprehensible; there would be no difficulty, for example, in programming a computer to react to stimuli along these lines (though this would not satisfy Descartes, for other reasons). Similarly, there is no difficulty in principle in programming a computer with a schematism that sharply restricts the form of a generative grammar, with an evaluation procedure for grammars of the given form, with a technique for determining whether given data are compatible with a grammar of the given form, with a fixed substructure of entities (such as distinctive features), rules, and principles, and so on – in short, with a universal grammar of the sort that has been proposed in recent years. For reasons that I have already mentioned, I believe that these proposals can be properly regarded as a further development of classical rationalist doctrine, as an elaboration of some of its main ideas regarding language and mind. Of course, such a theory will be “repugnant” to one who accepts empiricist doctrine and regards it as immune to question or challenge. It seems to me that this is the heart of the matter.

Putnam’s paper deals more directly with the points at issue, but it seems to me that his arguments are also inconclusive, because of certain incorrect assumptions that he makes about the nature of the acquired grammars. Putnam assumes that on the level of phonetics the only property proposed in universal grammar is that a language has “a short list

of phonemes.” This, he argues, is not a similarity among languages that requires elaborate explanatory hypotheses. The conclusion is correct; the assumption is quite wrong. In fact, as I have now pointed out several times, very strong empirical hypotheses have been proposed regarding the specific choice of universal features, conditions on the form and organisation of phonological rules, conditions on rule application, and so on. If these proposals are correct or near correct, then “similarities among languages” at the level of sound structure are indeed remarkable and cannot be accounted for simply by assumptions about memory capacity, as Putnam suggests.

Above the level of sound structure, Putnam assumes that the only significant properties of language are that they have proper names, that the grammar contains a phrase structure component, and that there are rules “abbreviating” sentences generated by the phrase structure component. He argues that the nature of the phrase structure component is determined by the existence of proper names; that the existence of a phrase structure component is explained by the fact that “all the natural measures of complexity of an algorithm – size of the machine table, length of computations, time, and space required for the computation – lead to the ... result”; that phrase structure systems provide the “algorithms which are ‘simplest’ for virtually any computing system,” hence also “for naturally evolved ‘computing systems’”; and that there is nothing surprising in the fact that languages contain rules of abbreviation.

Each of the three conclusions involves a false assumption. From the fact that a phrase structure system contains proper names one can conclude almost nothing about its other categories. In fact, there is much dispute at the moment about the general properties of the underlying phrase structure system for natural languages; the dispute is not in the least resolved by the existence of proper names.

As to the second point, it is simply untrue that all measures of complexity and speed of computation lead to phrase structure rules as the “simplest possible algorithm.” The only existing results that are even indirectly relevant show that context-free phrase structure grammars (a reasonable model for rules generating deep structures, when we exclude the lexical items and the distributional conditions they meet) receive an automata-theoretic interpretation as non-deterministic push-down storage automata, but the latter is hardly a “natural” notion from the point of view of “simplicity of algorithms” and so forth. In fact, it can be argued that the somewhat similar but not formally related concept of real-time deterministic automation is far more “natural” in terms of time and space conditions on computation.

However, it is pointless to pursue this topic, because what is at stake is not the “simplicity” of phrase structure grammars but rather of transformational grammars with a phrase structure component that plays a role in generating deep structures. And there is absolutely no mathematical concept of “ease of computation” or “simplicity of algorithm” that even vaguely suggests that such systems may have some advantage over the kinds of automata that have been seriously investigated from this point of view – for example, finite state automata, linear bounded automata, and so on. The basic concept of “structure-dependent operation” has never even been considered in a strictly mathematical concept. The source of this confusion is a misconception on Putnam’s part as to the nature of grammatical transformations. They are not rules that “abbreviate” sentences; rather, they are operations that form surface structures from underlying deep structures, in such ways as are illustrated in the preceding lecture and the references there cited. Hence, to show that transformational grammars are the “simplest possible” one would have to demonstrate that the “optimal” computing system would take a string of symbols as input and determine its surface structure, its

underlying deep structure, and the sequence of transformational operations that relates them. Nothing of the sort has been shown; in fact, the question has never even been raised.

Putnam argues that even if significant uniformities among languages were to be discovered, there would be a simpler explanation than the hypothesis of an innate universal grammar, namely their common origin. But this proposal involves a serious misunderstanding of the problem at issue. The grammar of a language must be discovered by the child from the data presented to him. As noted earlier, the empirical problem is to find a hypothesis about initial structure rich enough to account for the fact that a specific grammar is constructed by the child, but not so rich as to be falsified by the known diversity of language.

Questions of common origin are of potential relevance to this empirical issue in only one respect: If the existing languages are not a “fair sample” of the “possible languages,” we may be led mistakenly to propose too narrow a schema for universal grammar. However, as I mentioned earlier, the empirical problem that we face today is that no one has been able to devise an initial hypothesis rich enough to account for the acquisition by the child of the grammar that we are, apparently, led to attribute to him when we try to account for his ability to use the language in the normal way. The assumption of common origin contributes nothing to explaining how this achievement is possible. In short, the language is “reinvented” each time it is learned, and the empirical problem to be faced by the theory of learning is how this invention of grammar can take place.

Putnam does face this problem and suggests that there might be “general multipurpose learning strategies” that account for this achievement. It is, of course, an empirical question whether the properties of the “language faculty” are specific to language or are merely a

particular case of much more general mental faculties (or learning strategies).

This is a problem that has been discussed earlier in this lecture, inconclusively and in a slightly different context. Putnam takes for granted that it is only general “learning strategies” that are innate but suggests no grounds for this empirical assumption. As I have argued earlier, a non-dogmatic approach to this problem can be pursued, without reliance on unargued assumptions of this sort – that is, through the investigation of specific areas of human competence, such as language, followed by the attempt to devise a hypothesis that will account for the development of this competence. If we discover through such investigation that the same “learning strategies” are sufficient to account for the development of competence in various domains, we will have reason to believe that Putnam’s assumption is correct. If we discover that the postulated innate structures differ from case to case, the only rational conclusion would be that a model of mind must involve separate “faculties,” with unique or partially unique properties. I cannot see how anyone can resolutely insist on one or the other conclusion in the light of the evidence now available to us. But one thing is quite clear: Putnam has no justification for his final conclusion, that “invoking ‘Innateness’ only postpones the problem of learning; it does not solve it.” Invoking an innate representation of universal grammar does solve the problem of learning, if it is true that this is the basis for language acquisition, as it well may be. If, on the other hand, there are general learning strategies that account for the acquisition of grammatical knowledge, then postulation of an innate universal grammar will not “postpone” the problem of learning, but will rather offer an incorrect solution to this problem. The issue is an empirical one of truth or falsity, not a methodological one of states of investigation.

To summarise, it seems to me that neither Goodman nor Putnam offers a serious counter-argument to the proposals concerning innate mental structure that have been advanced (tentatively, of course, as befits empirical hypotheses) or suggests a plausible alternative approach, with empirical content, to the problem of acquisition of knowledge.

Assuming the rough accuracy of conclusions that seem tenable today, it is reasonable to suppose that a generative grammar is a system of many hundreds of rules of several different types, organised in accordance with certain fixed principles of ordering and applicability and containing a certain fixed substructure which, along with the general principles of organisation, is common to all languages. There is no *a priori* “naturalness” to such a system, any more than there is to the detailed structure of the visual cortex. No one who has given any serious thought to the problem of formalising inductive procedures or “heuristic methods” is likely to set much store by the hope that such a system as a generative grammar can be constructed by methods of any generality.

To my knowledge, the only substantive proposal to deal with the problem of acquisition of knowledge of language is the rationalist conception that I have outlined. To repeat: Suppose that we assign to the mind, as an innate property, the general theory of language that we have called “universal grammar.” This theory encompasses the principles that I discussed in the preceding lecture and many others of the same sort, and it specifies a certain subsystem of rules that provides a skeletal structure for any language and a variety of conditions, formal and substantive, that any further elaboration of the grammar must meet. The theory of universal grammar, then, provides a schema to which any particular grammar must conform. Suppose, furthermore, that we can make this schema sufficiently restrictive so that very few possible grammars conforming to the schema will be consistent with the meagre and degenerate data actually available to the language learner. His task, then, is to search among the possible

grammars and select one that is not definitely rejected by the data available to him. What faces the language learner, under these assumptions, is not the impossible task of inventing a highly abstract and intricately structured theory on the basis of degenerate data, but rather the much more manageable task of determining whether these data belong to one or another of a fairly restricted set of potential languages.

The tasks of the psychologist, then, divide into several sub-tasks. The first is to discover the innate schema that characterises the class of potential languages – that defines the “essence” of human language. This sub-task falls to that branch of human psychology known as linguistics; it is the problem of traditional universal grammar, of contemporary linguistic theory. The second sub-task is the detailed study of the actual character of the stimulation and the organism-environment interaction that sets the innate cognitive mechanism into operation. This is a study now being undertaken by a few psychologists, and it is particularly active right here in Berkeley. It has already led to interesting and suggestive conclusions. One might hope that such study will reveal a succession of maturational stages leading finally to a full generative grammar.

A third task is that of determining just what it means for a hypothesis about the generative grammar of a language to be “consistent” with the data of sense. Notice that it is a great oversimplification to suppose that a child must discover a generative grammar that accounts for all the linguistic data that has been presented to him and that “projects” such data to an infinite range of potential sound-meaning relations. In addition to achieving this, he must also differentiate the data of sense into those utterances that give direct evidence as to the character of the underlying grammar and those that must be rejected by the hypothesis he selects as ill-formed, deviant, fragmentary, and so on. Clearly, everyone succeeds in carrying out this task of differentiation – we all know, within tolerable limits of consistency, which sentences are well formed and literally

interpretable, and which must be interpreted as metaphorical, fragmentary, and deviant along many possible dimensions. I doubt that it has been fully appreciated to what extent this complicates the problem of accounting for language acquisition. Formally speaking, the learner must select a hypothesis regarding the language to which he is exposed that rejects a good part of the data on which this hypothesis must rest. Again, it is reasonable to suppose this is possible only if the range of tenable hypotheses is quite limited – if the innate schema of universal grammar is highly restrictive. The third sub-task, then, is to study what we might think of as the problem of “confirmation” – in this context, the problem of what relation must hold between a potential grammar and a set of data for this grammar to be confirmed as the actual theory of the language in question.

I have been describing the problem of acquisition of knowledge of language in terms that are more familiar in an epistemological than a psychological context, but I think that this is quite appropriate. Formally speaking, acquisition of “common-sense knowledge” – knowledge of a language, for example – is not unlike theory construction of the most abstract sort. Speculating about the future development of the subject, it seems to me not unlikely, for the reasons I have mentioned, that learning theory will progress by establishing the innately determined set of possible hypotheses, determining the conditions of interaction that lead the mind to put forth hypotheses from this set, and fixing the conditions under which such a hypothesis is confirmed – and, perhaps, under which much of the data is rejected as irrelevant for one reason or another.

Such a way of describing the situation should not be too surprising to those familiar with the history of psychology at Berkeley, where, after all, Edward Tolman has given his name to the psychology building; but I want to stress that the hypotheses I am discussing are qualitatively different in complexity and intricacy from anything that was considered in

the classical discussions of learning. As I have now emphasised several times, there seems to be little useful analogy between the theory of grammar that a person has internalised and that provides the basis for his normal, creative use of language, and any other cognitive system that has so far been isolated and described; similarly, there is little useful analogy between the schema of universal grammar that we must, I believe, assign to the mind as an innate character, and any other known system of mental organisation. It is quite possible that the lack of analogy testifies to our ignorance of other aspects of mental function, rather than to the absolute uniqueness of linguistic structure; but the fact is that we have, for the moment, no objective reason for supposing this to be true.

The way in which I have been describing acquisition of knowledge of language calls to mind a very interesting and rather neglected lecture given by [Charles Sanders Peirce](#) more than fifty years ago, in which he developed some rather similar notions about acquisition of knowledge in general. Peirce argued that the general limits of human intelligence are much more narrow than might be suggested by romantic assumptions about the limitless perfectibility of man (or, for that matter, than are suggested by his own “pragmaticist” conceptions of the course of scientific progress in his better-known philosophical studies). He held that innate limitations on admissible hypotheses are a precondition for successful theory construction, and that the “guessing instinct” that provides hypotheses makes use of inductive procedures only for “corrective action,” Peirce maintained in this lecture that the history of early science shows that something approximating a correct theory was discovered with remarkable ease and rapidity, on the basis of highly inadequate data, as soon as certain problems were faced; he noted “how few were the guesses that men of surpassing genius had to make before they rightly guessed the laws of nature.” And, he asked, “How was it that man was ever led to entertain that true theory? You cannot say that it

happened by chance, because the chances are too overwhelmingly against the single true theory in the twenty or thirty thousand years during which man has been a thinking animal, ever having come into any man's head." *A fortiori*, the chances are even more overwhelmingly against the true theory of each language ever having come into the head of every four-year-old child. Continuing with Peirce: "Man's mind has a natural adaptation to imagining correct theories of some kinds.... If man had not the gift of a mind adapted to his requirements, he could not have acquired any knowledge." Correspondingly, in our present case, it seems that knowledge of a language's grammar – can be acquired only by an organism that is "preset" with a severe restriction on the form of grammar. This innate restriction is a precondition, in the Kantian sense, for linguistic experience, and it appears to be the critical factor in determining the course and result of language learning. The child cannot know at birth which language he is to learn, but he must know that its grammar must be of a predetermined form that excludes many imaginable languages. Having selected a permissible hypothesis, he can use inductive evidence for corrective action, confirming or disconfirming his choice. Once the hypothesis is sufficiently well confirmed, the child knows the language defined by this hypothesis; consequently, his knowledge extends enormously beyond his experience and, in fact, leads him to characterise much of the data of experience as defective and deviant.

Peirce regarded inductive processes as rather marginal to the acquisition of knowledge; in his words, "Induction has no originality in it, but only tests a suggestion already made." To understand how knowledge is acquired, in the rationalist view that Peirce outlined, we must penetrate the mysteries of what he called "abduction," and we must discover that which "gives a rule to abduction and so puts a limit upon admissible hypotheses." Peirce maintained that the search for principles of abduction leads us to the study of innate ideas, which provide the instinctive

structure of human intelligence. But Peirce was no dualist in the Cartesian sense; he argued (not very persuasively, in my opinion) that there is a significant analogy between human intelligence, with its abductive restrictions, and animal instinct. Thus, he maintained that man discovered certain true theories only because his “instincts must have involved from the beginning certain tendencies to think truly” about certain specific matters; similarly, “You cannot seriously think that every little chicken that is hatched, has to rummage through all possible theories until it lights upon the good idea of picking up something and eating it. On the contrary, you think that the chicken has an innate idea of doing this; that is to say, that it can think of this, but has no faculty of thinking anything else.... But if you are going to think every poor chicken endowed with an innate tendency towards a positive truth, why should you think to man alone this gift is denied?”

No one took up Peirce’s challenge to develop a theory of abduction, to determine those principles that limit the admissible hypotheses or present them in a certain order. Even today, this remains a task for the future. It is a task that need not be undertaken if empiricist psychological doctrine can be substantiated; therefore, it is of great importance to subject this doctrine to rational analysis, as has been done, in part, in the study of language. I would like to repeat that it was the great merit of structural linguistics, as of Hullian learning theory in its early stages and of several other modern developments, to have given precise form to certain empiricist assumptions.” Where this step has been taken, the inadequacy of the postulated mechanisms has been clearly demonstrated, and, in the case of language at least, we can even begin to see just why any methods of this sort must fail – for example, because they cannot, in principle, provide for the properties of deep structures and the abstract operations of formal grammar. Speculating about the future, I think it is not unlikely that the dogmatic character of the general empiricist framework and its

inadequacy to human and animal intelligence will gradually become more evident as specific realisations, such as taxonomic linguistics, behaviourist learning theory, and the perception models,” heuristic methods, and “general problem solvers” of the early enthusiasts of “artificial intelligence,” are successively rejected on empirical grounds when they are made precise and on grounds of vacuity when they are left vague. And – assuming this projection to be accurate – it will then be possible to undertake a general study of the limits and capacities of human intelligence, to develop a Peircean logic of abduction.

Modern psychology is not devoid of such initiatives. The contemporary study of generative grammar and its universal substructure and governing principles is one such manifestation. Closely related is the study of the biological bases of human language, an investigation to which Eric Lenneberg has made substantial contributions.” It is tempting to see a parallel development in the very important work of Piaget and others interested in “genetic epistemology,” but I am not sure that this is accurate. It is not clear to me, for example, what Piaget takes to be the basis for the transition from one of the stages that he discusses to the next, higher stage. There is, furthermore, a possibility, suggested by recent work of Mehler and Bever,” that the deservedly well-known results on conservation, in particular, may not demonstrate successive stages of intellectual development in the sense discussed by Piaget and his coworkers, but something rather different. If the preliminary results of Mehler and Bever are correct, then it would follow that the “final stage,” in which conservation is properly understood, was already realised at a very early period of development. Later, the child develops a heuristic technique that is largely adequate but that fails under the conditions of the conservation experiment. Still later, he adjusts this technique successfully and once again makes the correct judgments in the conservation experiment. If this analysis is correct, then what we are observing is not a

succession of stages of intellectual development, in Piaget's sense, but rather slow progress in bringing heuristic techniques into line with general concepts that have always been present. These are interesting alternatives; either way, the results may bear in important ways on the topics we are considering.

Still more clearly to the point, I think, are the developments in comparative ethology over the past thirty years, and certain current work in experimental and physiological psychology. One can cite many examples: for example, in the latter category, the work of Bower, suggesting an innate basis for the perceptual constancies; studies in the Wisconsin primate laboratory on complex innate releasing mechanisms in rhesus monkeys; the work of Hubel, Barlow, and others on highly specific analysing mechanisms in the lower cortical centers of mammals; and a number of comparable studies of lower organisms (for example, the beautiful work of Lettvin and his associates on frog vision). There is now good evidence from such investigations that perception of line, angle, motion, and other complex properties of the physical world is based on innate organisation of the neural system.

In some cases at least, these built-in structures will degenerate unless appropriate stimulation takes place at an early stage in life, but although such experience is necessary to permit the innate mechanisms to function, there is no reason to believe that it has more than a marginal effect on determining how they function to organise experience. Furthermore, there is nothing to suggest that what has so far been discovered is anywhere near the limit of complexity of innate structures. The basic techniques for exploring the neural mechanisms are only a few years old, and it is impossible to predict what order of specificity and complexity will be demonstrated when they come to be extensively applied. For the present, it seems that most complex organisms have highly specific forms of sensory and perceptual organisation that are associated with the *Umwelt*

and the manner of life of the organism. There is little reason to doubt that what is true of lower organisms is true of humans as well. Particularly in the case of language, it is natural to expect a close relation between innate properties of the mind and features of linguistic structure; for language, after all, has no existence apart from its mental representation. Whatever properties it has must be those that are given to it by the innate mental processes of the organism that has invented it and that invents it anew with each succeeding generation, along with whatever properties are associated with the conditions of its use. Once again, it seems that language should be, for this reason, a most illuminating probe with which to explore the organisation of mental processes.

Turning to comparative ethology, it is interesting to note that one of its earliest motivations was the hope that through the “investigation of the *a priori*, of the innate working hypotheses present in subhuman organisms,” it would be possible to shed light on the *a priori* forms of human thought. This formulation of intent is quoted from an early and little-known paper by Konrad Lorenz.” Lorenz goes on to express views very much like those Peirce had expressed a generation earlier. He maintains:

One familiar with the innate modes of reaction of subhuman organisms can readily hypothesise that the *a priori* is due to hereditary differentiations of the central nervous system which have become characteristic of the species, producing hereditary dispositions to think in certain forms.... Most certainly Hume was wrong when he wanted to derive all that is *a priori* from that which the senses supply to experience, just as wrong as Wundt or Helmholtz who simply explain it as an abstraction from preceding experience. Adaptation of the *a priori* to the real world has no more originated from “experience” than adaptation of the fin of the fish to the properties of water. just as the form of the fin is given *a priori*, prior to any individual negotiation of the young fish with the water, and just as it is this form that makes possible this negotiation, so it is also the case with our forms of perception and categories in their relationship to our negotiation with

the real external world through experience. In the case of animals, we find limitations specific to the forms of experience possible for them. We believe we can demonstrate the closest functional and probably genetic relationship between these animal *a priori*'s and our human *a priori*. Contrary to Hume, we believe, just as did Kant, that a "pure" science of innate forms of human thought, independent of all experience, is possible.

Peirce, to my knowledge, is original and unique in stressing the problem of studying the rules that limit the class of possible theories. Of course, his concept of abduction, like Lorenz's biological *a priori*, has a strongly Kantian flavour, and all derive from the rationalist psychology that concerned itself with the forms, the limits, and the principles that provide "the sinews and connections" for human thought, that underlie "that infinite amount of knowledge of which we are not always conscious," of which Leibnitz spoke. It is therefore quite natural that we should link these developments to the revival of philosophical grammar, which grew from the same soil as an attempt, quite fruitful and legitimate, to explore one basic facet of human intelligence.

In recent discussion, models and observations derived from ethology have frequently been cited as providing biological support, or at least analogue, to new approaches to the study of human intelligence. I cite these comments of Lorenz's mainly in order to show that this reference does not distort the outlook of at least some of the founders of this domain of comparative psychology.

One word of caution is necessary in referring to Lorenz, now that he has been discovered by Robert Ardrey and Joseph Alsop and popularised as a prophet of doom. It seems to me that Lorenz's views on human aggression have been extended to near absurdity by some of his expositors. It is no doubt true that there are innate tendencies in the human psychic constitution that lead to aggressiveness under specific social and

cultural conditions. But there is little reason to suppose that these tendencies are so dominant as to leave us forever tottering on the brink of a Hobbesian war of all against all – as, incidentally, Lorenz at least is fully aware, if I read him rightly. Scepticism is certainly in order when a doctrine of man's "inherent aggressiveness" comes to the surface in a society that glorifies competitiveness, in a civilisation that has been distinguished by the brutality of the attack that it has mounted against less fortunate peoples. It is fair to ask to what extent the enthusiasm for this curious view of man's nature is attributable to fact and logic and to what extent it merely reflects the limited extent to which the general cultural level has advanced since the days when Clive and the Portuguese explorers taught the meaning of true savagery to the inferior races that stood in their way.

In any event, I would not want what I am saying to be confused with other, entirely different attempts to revive a theory of human instinct. What seems to me important in ethology is its attempt to explore the innate properties that determine how knowledge is acquired and the character of this knowledge. Returning to this theme, we must consider a further question: How did the human mind come to acquire the innate structure that we are led to attribute to it? Not too surprisingly, Lorenz takes the position that this is simply a matter of natural selection. Peirce offers a rather different speculation, arguing that "nature fecundates the mind of man with ideas which, when these ideas grow up, will resemble their father, Nature." Man is "provided with certain natural beliefs that are true" because "certain uniformities ... prevail throughout the universe, and the reasoning mind is [it]self a product of this universe. These same laws are thus, by logical necessity, incorporated in his own being." Here, it seems clear that Peirce's argument is entirely without force and that it offers little improvement over the pre-established harmony that it was presumably intended to replace. The fact that the mind is a product of

natural laws does not imply that it is equipped to understand these laws or to arrive at them by “abduction.” There would be no difficulty in designing a device (say, programming a computer) that is a product of natural law, but that, given data, will arrive at any arbitrary absurd theory to “explain” these data.

In fact, the processes by which the human mind achieved its present stage of complexity and its particular form of innate organisation are a total mystery, as much so as the analogous questions about the physical or mental organisation of any other complex organism. It is perfectly safe to attribute this development to “natural selection,” so long as we realise that there is no substance to this assertion, that it amounts to nothing more than a belief that there is some naturalistic explanation for these phenomena. The problem of accounting for evolutionary development is, in some ways, rather like that of explaining successful abduction. The laws that determine possible successful mutation and the nature of complex organisms are as unknown as the laws that determine the choice of hypotheses.” With no knowledge of the laws that determine the organisation and structure of complex biological systems, it is just as senseless to ask what the “probability” is for the human mind to have reached its present state as it is to inquire into the “probability” that a particular physical theory will be devised. And, as we have noted, it is idle to speculate about laws of learning until we have some indication of what kind of knowledge is attainable – in the case of language, some indication of the constraints on the set of potential grammars.

In studying the evolution of mind, we cannot guess to what extent there are physically possible alternatives to, say, transformational generative grammar, for an organism meeting certain other physical conditions characteristic of humans. Conceivably, there are none – or very few in which case talk about evolution of the language capacity is beside the point. The vacuity of such speculation, however, has no bearing one way

or another on those aspects of the problem of mind that can be sensibly pursued. It seems to me that these aspects are, for the moment, the problems illustrated in the case of language by the study of the nature, the use, and the acquisition of linguistic competence.

There is one final issue that deserves a word of comment. I have been using mentalistic terminology quite freely, but entirely without prejudice as to the question of what may be the physical realisation of the abstract mechanisms postulated to account for the phenomena of behaviour or the acquisition of knowledge. We are not constrained, as was Descartes, to postulate a second substance when we deal with phenomena that are not expressible in terms of matter in motion, in his sense. Nor is there much point in pursuing the question of psychophysical parallelism, in this connection. It is an interesting question whether the functioning and evolution of human mentality can be accommodated within the framework of physical explanation, as presently conceived, or whether there are new principles, now unknown, that must be invoked, perhaps principles that emerge only at higher levels of organisation than can now be submitted to physical investigation. We can, however, be fairly sure that there will be a physical explanation for the phenomena in question, if they can be explained at all, for an uninteresting terminological reason, namely that the concept of "physical explanation" will no doubt be extended to incorporate whatever is discovered in this domain, exactly as it was extended to accommodate gravitational and electromagnetic force, massless particles, and numerous other entities and processes that would have offended the common sense of earlier generations. But it seems clear that this issue need not delay the study of the topics that are now open to investigation, and it seems futile to speculate about matters so remote from present understanding.

I have tried to suggest that the study of language may very well, as was traditionally supposed, provide a remarkably favourable perspective for

the study of human mental processes. The creative aspect of language use, when investigated with care and respect for the facts, shows that current notions of habit and generalisation, as determinants of behaviour or knowledge, are quite inadequate. The abstractness of linguistic structure reinforces this conclusion, and it suggests further that in both perception and learning the mind plays an active role in determining the character of the acquired knowledge. The empirical study of linguistic universals has led to the formulation of highly restrictive and, I believe, quite plausible hypotheses concerning the possible variety of human languages, hypotheses that contribute to the attempt to develop a theory of acquisition of knowledge that gives due place to intrinsic mental activity. It seems to me, then, that the study of language should occupy a central place in general psychology.

Surely the classical questions of language and mind receive no final solution, or even the hint of a final solution, from the work that is being actively pursued today. Nevertheless, these problems can be formulated in new ways and seen in a new light. For the first time in many years, it seems to me, there is some real opportunity for substantial progress in the study of the contribution of the mind to perception and the innate basis for acquisition of knowledge. Still, in many respects, we have not made the first approach to a real answer to the classical problems. For example, the central problems relating to the creative aspect of language use remain as inaccessible as they have always been. And the study of universal semantics, surely crucial to the full investigation of language structure, has barely advanced since the medieval period. Many other critical areas might be mentioned where progress has been slow or nonexistent. Real progress has been made in the study of the mechanisms of language, the formal principles that make possible the creative aspect of language use and that determine the phonetic form and semantic content of utterances. Our understanding of these mechanisms, though only fragmentary, does

seem to me to have real implications for the study of human psychology. By pursuing the kinds of research that now seem feasible and by focusing attention on certain problems that are now accessible to study, we may be able to spell out in some detail the elaborate and abstract computations that determine, in part, the nature of percepts and the character of the knowledge that we can acquire the highly specific ways of interpreting phenomena that are, in large measure, beyond our consciousness and control and that may be unique to man.