BLACKWELL PHILOSOPHY ANTHOLOGIES

Each volume in this outstanding series provides an authoritative and comprehensive collection of the essential primary readings from philosophy's main fields of study. Designed to complement the *Blackwell Companions to Philosophy* series, each volume represents an unparalleled resource in its own right, and will provide the ideal platform for course use.

1 Cottingham: *Western Philosophy: An Anthology*
2 Cahoone: *From Modernism to Postmodernism: An Anthology*
3 LaFollette: *Ethics in Practice: An Anthology*
4 Goodin and Pettit: *Contemporary Political Philosophy: An Anthology*
5 Eze: *African Philosophy: An Anthology*
6 McNeill and Feldman: *Continental Philosophy: An Anthology*
7 Kim and Sosa: *Metaphysics: An Anthology*
8 Lycan: *Mind and Cognition: An Anthology* (second edition)
9 Kuhse and Singer: *Bioethics: An Anthology*
10 Cummins and Cummins: *Minds, Brains, and Computers: The Foundations of Cognitive Science: An Anthology*
11 Sosa and Kim: *Epistemology: An Anthology*

Minds, Brains, and Computers
The Foundations of Cognitive Science
An Anthology

Edited by

Robert Cummins
University of California at Davis

and

Denise Dellarosa Cummins
University of California at Davis

Chomsky, "Linguistics & Philosophy" 4

© 2000
BLACKWELL
Publishers

30

Noam Chomsky

The methods and concerns of linguists and philosophers are similar in so many respects that it would be folly, I believe, to insist on a sharp separation of these disciplines, or for either to maintain a parochial disregard for insights achieved in the other. A number of examples might be cited to illustrate the possibility of fruitful interchange between the two. Zeno Vendler, in his recent book *Linguistics and Philosophy*, goes so far as to maintain that "the science of structural linguistics" provides "a new technique" for analytic philosophy, one that "is nothing but the natural continuation of the line of development that goes through the philosophers of ordinary language to J. L. Austin." For reasons to which I will return in a moment, I am a bit skeptical about the contribution that linguistics might make to philosophy along the lines that he sketches, but I think that he has shown that certain concepts of linguistics can be used in a rewarding way in the investigation of problems that have arisen in analytic philosophy.

Conversely, as the attention of linguists begins to turn to problems of meaning and use, there is no question that they can learn much from the long tradition of philosophical investigation of such problems, although here too, I think, a note of skepticism is in order.

To facilitate the discussion of this and other topics, let me present a small illustration of a problem that is at the frontier of research today. In the descriptive study of any language a central problem is to formulate a set of rules that generate what we may call the "surface structures" of utterances. By the term "surface structure," I refer to the analysis of an utterance into a hierarchy of phrases, each belonging to a specific category. This hierarchy can be represented as a labeled bracketing of the utterance, in an obvious sense. For example, consider the two sentences:

(1) John is certain that Bill will leave
(2) John is certain to leave

The surface structures of these utterances can be represented, in a natural way, with the following labeled bracketing:

(1′) $[_S[_{NP}\text{ John}]\ [_{VP}\text{ is }[_{AP}\text{ certain}$
$[_S\text{ that }[_{NP}\text{ Bill}]_{VP}\text{ will leave}]]]]$
(2′) $[_S[_{NP}\text{ John}]\ [_{VP}\text{ is }[_{AP}\text{ certain}]\ [_{VP}\text{ to leave}]]]$

Paired brackets bound phrases; the label assigned to a pair of brackets indicates the category of the bounded phrase. Thus in (1), "certain that Bill will leave" is a phrase of the category, Adjective Phrase; in both (1) and (2), "John" is a phrase of the category, Noun Phrase; "will leave" is a Verb Phrase in (1); and both (1) and (2) are phrases of the category, Sentence. One may question the details of these particular analyses, but there is little doubt that at some level of description, these, or representations very much like them, constitute a significant aspect of the structure of the sentences (1) and (2), and, more generally, that every sentence of the language has a surface structure of roughly this sort. There is, for example, strong evidence that the perceived phonetic form of the utterance is determined, by phonological rules of considerable generality, from representations of essentially this sort.

Granting this much, the linguist studying English will try to formulate a set of rules that generate an infinite number of surface structures, one for each sentence of English. Correspondingly, linguistic theory will be concerned with the problem of how such structures are generated in any human language, and will try to formulate general principles governing the systems of rules that express the facts of one or another such language.

Given the evidence available to us today, it seems to me reasonable to propose that in every human language surface structures are generated from structures of a more abstract sort, which I will refer to as "deep structures," by certain formal operations of a very special kind generally called "grammatical transformations." Each transformation is a mapping of labeled bracketings onto labeled bracketings. Deep structures are themselves labeled bracketings. The infinite class of deep structures is specified by a set of "base rules." Transformations applied in sequence to deep structures in accordance with certain fixed conventions and principles ultimately generate the surface structures of the sentences of the language. Thus a set of base rules defining an infinite class of deep structures and a set of grammatical transformations can serve to generate the surface structures.

To illustrate, consider sentences (1) and (2) again. The underlying deep structures might be represented roughly in the form (1″), (2″):

(1″) same as (1′)
(2″) $[_S[_{NP}[_S[_{NP}\text{ John}]$
$[_{VP}\ \textit{to}\text{ leave}]]]\ [_{VP}\ \textit{is}\ [_{AP}\text{ certain}]]]$

We may think of these deep structures as expressing the fact that in (1), we predicate of John that he is certain that Bill will leave, whereas in (2), which is rather similar to (1) in surface structure, we predicate of the proposition that John leaves, that it is certain, in a very different sense of "certain." There is no difficulty in defining the concepts Subject and Predicate, in terms of configurations in deep structures, so that they express the intended interpretation. The operations that derive (2′) from (2″) include an operation of "extraposition," which from a structure very much like (2″) would yield the structure (3), and an operation of "it-replacement" which derives (2′) from a structure almost exactly like (3), but with "to" in place of "will" and "that" deleted:

(3) $[_S[_{NP}\text{ it}]\ [_{VP}\text{ is }[_{AP}\text{ certain}]$
$[_S\text{ that }[_{NP}\text{ John}]\ [_{VP}\text{ will leave}]]]]$

Details aside, the theory of "transformational generative grammar" maintains that all surface structures are formed by application of such transformations – each of which maps labeled bracketings onto labeled bracketings – from deep structures that are often quite abstract. The sentences (1) and (2) are similar in surface structure, but very different in deep structure; the sentences (2) and (3) are very similar in deep structure, but quite different in surface structure. The deep structures of the language are quite restricted in their variety, and it appears that there are universal conditions that sharply restrict the class of possible rules.

Consider now the matter of semantic interpretation. It is clear from these quite typical examples that the surface structures give little indication of the semantic interpretation, whereas the deep structures are quite revealing, in this respect. Pursuing this line of reasoning, one might propose a further elaboration of the theory just outlined, in the following terms. Let us suppose that there is a system of "universal semantics" that specifies the class of possible semantic representations for a natural language much in the way that universal phonetics specifies the class of possible phonetic representations, by specifying a class of distinctive features and certain conditions on their combination. Observe that it would be perfectly reasonable to study universal semantics even without any clear idea as to what its constituent elements might be, just as one could draw fairly persuasive conclusions regarding universal phonetics from consideration of the slow growth of the number of distinct sentences with increasing length, the phenomena of rhyme and assonance, the lack of slow drift through the "space" of sentences under chains of repetition, etc., even without any conception of what the distinctive features of this system might be. In any event, still supposing this to be a reasonable approach, one might propose that a language contains rules associating

Chomsky, N., "Linguistics and philosophy," from S. Hook (ed.), *Language and Philosophy* (New York University Press, New York, 1969).

deep structures with representations drawn from universal semantics, as it contains phonological rules relating surface structures to representations drawn from universal phonetics.

At this point in the development of such a theory, the linguist would do well to turn to work in analytic philosophy, particularly to the many studies of referential opacity. One essential empirical assumption in the preceding account is that surface structure cannot contribute to meaning; whatever contribution the expression P makes to the meaning of the sentence XPY must be determined by the deep structure underlying P. The investigation of referential opacity has turned up a great number of examples illustrating how replacement of one expression by another changes meaning, even when the semantic connection between the two is very close. The approach just outlined would have to guarantee that in each such case there is a corresponding difference in deep structure to which the difference in meaning can be attributed. Without pursuing the matter, I would simply note that the nature of these examples makes it appear very unlikely that such an approach can succeed; but, in any event, the study of this aspect of linguistic theory must certainly take into account a mass of evidence that has been accumulated in the course of philosophical investigation.

I have mentioned the possibility that insights developed in the course of philosophical analysis might be relevant to the study of a central part of linguistic theory, and that concepts of linguistics might be useful to the philosopher in his work. Nevertheless, it seems to me that one should not expect too much from an interchange of this sort, for a number of reasons. In the cases I have mentioned, what is proposed is that the incidental by-products of research in one field will be of use for the central concerns of another. Furthermore, it is a fact that neither field makes use of research techniques of a sophisticated or specialized nature. Thus one would expect that in each field, it would be quite possible to collect and analyze the information relevant to its specific concerns directly. It is, therefore, something of an accident when one field can build directly on results of the other.

For these reasons, I think that Vendler may be expecting too much of the method he suggests, namely, "an appeal to the facts of language already organized by the science of structural linguistics." I believe that modern linguistics has real achievements to its credit, and that some of these do have relevance to philosophical questions. But it must be kept in mind that these achievements owe little to modern science and less to modern technology. The gathering of data is informal; there has been very little use of experimental approaches (outside of phonetics) or of complex techniques of data collection and data analysis of a sort that can easily be devised, and that are widely used in the behavioral sciences. The arguments in favor of this informal procedure seem to me quite compelling; basically, they turn on the realization that for the theoretical problems that seem most critical today, it is not at all difficult to obtain a mass of crucial data without use of such techniques. Consequently, linguistic work, at what I believe to be its best, lacks many of the features of the behavioral sciences. Nor is it obvious that the development of explanatory theories in linguistics merits the honorific designation "scientific." I think that these intellectual constructions are nontrivial and often illuminating. However, apart from certain insights owed to modern logic and mathematics, there is no reason why they could not have been developed many years ago. In fact, were it not for the dominance of certain empiricist assumptions to which I will return directly, I suspect that they would have been developed long before now and that much of what is new and exciting in linguistics today would be taken for granted by any educated person.

There are many questions about language that a philosopher might ask to which linguistics provides no answer and no reasonable hope for an answer. For example, a philosopher concerned with problems of knowledge, or causality (to take an example of Vendler's), might well be interested in investigating in detail the properties of the words "know" and "cause." Since linguistics offers no privileged access to data of this sort, it would be merely a lucky accident if acquaintance with linguistics proved to be of substantial help in this inquiry. A linguistic form is not of importance to linguistics because of the intrinsic interest of the concept or proposition it expresses (if any), but because of the evidence it provides concerning some assumption about the nature of language. Thus the analysis of sentences (1), (2), and (3) has been of interest to linguistics because of the light it sheds on the nature of deep and surface structures and the grammatical transformations that link them. Such data are of importance to linguistics insofar as they can be explained on the basis of some interesting assumptions about the

organization of grammar, and are inconsistent with other such assumptions. In themselves, these facts are of no more interest than the fact that certain marks appear on a photographic plate at the base of a South African mine shaft. The latter is critical for elementary particle theory for the same reason that the facts related to sentences (1)–(3) are important for the theory of language. Similar remarks can be made about the likelihood that the conclusions of philosophers or the data they accumulate will be important for linguistics.

To make the matter more concrete, consider again the examples (1)–(3). Conceivably, such sentences, and others like them, might be of some interest to a philosopher concerned with the various concepts of certainty. These examples are of interest to linguistics, at the moment, for entirely different reasons. Thus it is interesting that there is a nominalized expression corresponding to (1), but no nominalized expression corresponding to (2); (4) is a nominalized form of (1), but we cannot form (5), corresponding to (2):

(4) John's certainty that Bill would leave
(5) John's certainty to leave

The distinction is more general; thus consider (6) and (7):

(6) John is eager to leave
(7) John is easy to leave

Corresponding to (6), we have the nominal phrase (8); but we cannot form (9) corresponding to (7):

(8) John's eagerness to leave
(9) John's easiness to leave

Notice that sentence (6) is like (1) in that the deep structure is very close to the surface structure; whereas (7) is like (2) in that the deep structure is very different from the surface structure. In fact, the surface structure of (7) would be formed by operations much like those that form (2) from (2″) and (3), by a derivation of roughly the form (10):

(10)(i) $[_S$ for one to leave John$]_S$ is easy (analogous to (2″))
(ii) it is easy $[_S$ for one to leave John$]_S$ (analogous to (3))
(iii) John is easy to leave (= (7), analogous to (2))

The generalization exemplified by (1), (2), (4)–(9) is that a nominal phrase can be formed corresponding to a base structure but not to a surface structure. Thus we have (4) corresponding to (1″) and (8) corresponding to (6) (more properly, to the deep structure underlying (6) as (1′) underlies (1)), but no nominalized expression such as (5) and (9), corresponding to the surface structures (2) and (7). This general observation can be illustrated by many other examples. It is interesting because of the support it lends to the assumption that abstract deep structures of the sort illustrated play a role in the mental representation of sentences. We find that when we study English grammar on the basis of this and related assumptions, we are able to characterize quite readily the class of sentences to which there correspond nominal phrases of the sort under discussion. There is no natural way to characterize this class in terms of surface structure, since, as we have seen, sentences that are very similar in surface structure behave quite differently with respect to the formal processes involved in the construction of nominal expressions. We might go on to try to explain these facts at a deeper level by formulating a principle of universal grammar from which it would follow that the nominal phrases in question will correspond only to deep structures.

To summarize, the examples in question are important for the study of language because of the evidence that they provide in support of a particular theory of linguistic structure, not because of the fact that the various concepts of certainty are of interest in their own right. The philosopher concerned with certainty would learn very little from a collection of data that is of great interest for linguistic research.

Apart from accident or matters of personal history, linguistics will be of relevance to philosophy only insofar as its conclusions about the nature of language bear on questions that concern the philosopher. One cannot predict to what extent this will be true in the future; it might turn out, for example, that linguistic study of semantic and syntactic structure in the future will provide a firm basis for certain kinds of philosophical investigation – one thinks, for example, of the potential relevance of a systematic classification of verbs that would have cross-language validity. For the moment, this is more a hope for the future than a present reality, however. Still, I think that a case can be made that certain well-founded conclusions about the nature of language do bear on traditional philosophical questions, but in ways rather different from those just mentioned. Specifically, I think that these conclusions are relevant to the problem of how knowledge is acquired and how the character of human knowledge is determined by certain general properties of the mind. What I would like to do, for the remainder of this paper, is to restate certain proposals about this matter that have been developed elsewhere,[1] and then to consider a variety of problems and objections that have been raised by several philosophers with respect to these proposals.[2]

One might adopt the following research strategy for the study of cognitive processes in humans. A person is presented with a physical stimulus that he interprets in a certain way. Let us say that he constructs a certain "precept" that represents certain of his conclusions (unconscious, in general) about the source of stimulation. To the extent that we can characterize this precept, we can proceed to investigate the process of interpretation. We can, in other words, proceed to develop a model of perception that takes stimuli as inputs and assigns precepts as "outputs," a model that will meet certain given empirical conditions on the actual pairing of stimuli with interpretations of these stimuli. For example, the person who understands sentences (1) and (2) knows (whether he is aware of it or not) that in the case of (2) it is a proposition that is certain and in the case of (1) it is a person who is certain of something, in a very different sense of "certain." If we are interested in studying perception of language, specifically, the processes by which sentences are understood, we can begin by describing the percepts in such a way as to bring out this difference, as we did in proposing that (1″) and (2″), interpreted in the suggested manner, are essential components of the percept. We can then ask how these percepts are constructed by the hearer, given the input stimuli (1) and (2).

A perceptual model that relates stimulus and percept might incorporate a certain system of beliefs, certain strategies that are used in interpreting stimuli, and other factors – for example, organization of memory. In the case of language, the technical term for the underlying system of beliefs is "grammar," or "generative grammar." A grammar is a system of rules that generates an infinite class of "potential percepts," each with its phonetic, semantic, and syntactic aspects, the class of structures that constitutes the language in question. The percepts themselves are first-order constructs; we determine their properties by experiment and observation. The grammar that underlies the formation of percepts is a second-order construct. To study it, we must abstract away from the other factors that are involved in the use and understanding of language, and concentrate on the knowledge of language[3] that has been internalized in some manner by the language user.

Concentrating on this system, we can then inquire into the means by which it was acquired and the basis for its acquisition. We can, in other words, attempt to construct a second model, a learning model, which takes certain data as input and gives, as "output," the system of beliefs that is one part of the internal structure of the perceptual model. The "output," in this case, is represented in the "final state" of the organism that has acquired this system of beliefs; we are asking, then, how this final state was achieved, through the interplay of innate factors, maturational processes, and organism–environment interaction.

In short, we can begin by asking "what is perceived" and move from there to a study of perception. Focusing on the role of belief (in our case, knowledge of language) in perception, we can try to characterize "what is learned" and move from there to the study of learning. One might, of course, decide to study some other topic, or to proceed in some different manner. Thus much of modern psychology has decided, for reasons that do not impress me, to limit itself to the study of behavior and control of behavior. I do not want to pursue the matter here, but I will merely state my own opinion: that this approach has proven quite barren, and that it is irrational to limit one's objectives in this way. One cannot hope to study learning or perception in any useful way by adhering to methodological strictures that limit the conceptual apparatus so narrowly as to disallow the concept "what is perceived" and the concept "what is learned."

I think that interesting conclusions can be reached when one studies human language along the lines just outlined. In the areas of syntax and phonetics at least, a plausible general account can be given of the system of representation for percepts in any human language. Furthermore, there has been substantial progress in constructing generative grammars that express the knowledge of language that is the "output" of a learning model and a fundamental component of a perceptual model. There is, I believe, good evidence that a

generative grammar for a human language contains a system of base rules of a highly restricted sort, a set of grammatical transformations that map the deep structures formed in accordance with base rules onto surface structures, and a set of phonological rules that assign phonetic interpretations, in a universal phonetic alphabet, to surface structures. Furthermore, there is also good evidence that certain highly restrictive principles determine the functioning of these rules, conditions of ordering and organization of a complex and intricate sort. There is a considerable literature dealing with these matters, and I will not try to review it here. I only wish to emphasize that there is no a priori necessity for a language to be organized in the highly specific manner proposed in these investigations. Hence if this theory of linguistic structure is correct, or near correct, some nontrivial problems arise for the theory of human learning. Specifically, we must ask how on the basis of the limited data available to him, the child is able to construct a grammar of the sort that we are led to ascribe to him, with its particular choice and arrangement of rules and with the restrictive principles of application of such rules. What, in other words, must be the internal structure of a learning model that can duplicate this achievement? Evidently, we must try to characterize innate structure in such a way as to meet two kinds of empirical conditions. First, we must attribute to the organism, as an innate property, a structure rich enough to account for the fact that the postulated grammar is acquired on the basis of the given conditions of access to data; second, we must not attribute to the organism a structure so rich as to be incompatible with the known diversity of languages. We cannot attribute knowledge of English to the child as an innate property, because we know that he can learn Japanese as well as English. We cannot attribute to him merely the ability to form associations, or to apply the analytic procedures of structural linguistics, because (as is easy to show when these proposals are made precise) the structures they yield are not those that we must postulate as generative grammars. Within the empirical bounds just stated, we are free to construct theories of innate structure and to test them in terms of their empirical consequences. To say this is merely to define the problem. Substantive questions arise only when a specific theory is proposed.

By investigating sentences and their structural descriptions, speech signals and the percepts to which they give rise, we can arrive at detailed conclusions regarding the generative grammar that is one fundamental element in linguistic performance, in speech and understanding of speech. Turning then to the next higher level of abstraction, we raise the question of how this generative grammar is acquired. From a formal point of view, the grammar that is internalized by every normal human can be described as a theory of his language, a theory of a highly intricate and abstract form that determines, ultimately, a connection between sound and meaning by generating structural descriptions of sentences ("potential percepts"), each with its phonetic, semantic, and syntactic aspects. From this point of view, one can describe the child's acquisition of knowledge of language as a kind of theory construction. Presented with highly restricted data, he constructs a theory of the language of which this data is a sample (and, in fact, a highly degenerate sample, in the sense that much of it must be excluded as irrelevant and incorrect – thus the child learns rules of grammar that identify much of what he has heard as ill-formed, inaccurate and inappropriate). The child's ultimate knowledge of language obviously extends far beyond the data presented to him. In other words, the theory he has in some way developed has a predictive scope of which the data on which it is based constitute a negligible part. The normal use of language characteristically involves new sentences, sentences that bear no point-by-point resemblance or analogy to those in the child's experience. Furthermore, the task of constructing this system is carried out in a remarkably similar way by all normal language learners, despite wide differences in experience and ability. The theory of human learning must face these facts.

I think that these facts suggest a theory of human intelligence that has a distinctly rationalist flavor. Using terms suggested by Peirce, in his lectures on "the logic of abduction," the problem of the theory of learning is to state the condition that "gives a rule to abduction and so puts a limit on admissible hypotheses." If "man's mind has a natural adaptation to imagining correct theories of some kinds," then acquisition of knowledge of a sort that we are considering is possible. The problem for the psychologist (or linguist) is to formulate the principles that set a limit to admissible hypotheses. I have made detailed suggestions in this regard elsewhere, and will not repeat them here. Roughly, I think it reasonable to postulate that the principles of general linguistics regarding the nature of rules, their organization, the principles by which they function, the kinds or representations to which they apply and which they form, all constitute part of the innate condition that "puts a limit on admissible hypotheses." If this suggestion is correct, then there is no more point asking how these principles are learned than there is in asking how a child learns to breathe, or, for that matter, to have two arms. Rather, the theory of learning should try to characterize the particular strategies that a child uses to determine that the language he is facing is one, rather than another, of the "admissible languages." When the principles just alluded to are made precise, they constitute an empirical assumption about the innate basis for the acquisition of knowledge, an assumption that can be tested in a variety of ways. In particular, we can ask whether it falls between the bounds described earlier: that is, does it ascribe a rich enough innate structure to account for the acquisition of knowledge, but a structure not so rich as to be falsified by the diversity of languages. We might also ask many other questions, for example, how the schema that is proposed as a basis for acquisition of knowledge of language relates to the principles that "give a rule to abduction" in other domains of human (or animal) intelligence.

What I am suggesting is that if we wish to determine the relevance of linguistics to philosophy, we must investigate the conclusions that can be established concerning the nature of language, the ways in which language is used and understood, the basis for its acquisition. I think that these conclusions have interesting consequences for psychological theory, in particular, that they strongly support an account of mental processes that is, in part, familiar, from rationalist speculation about these matters. They support the conclusion that the role of intrinsic organization is very great in perception, and that a highly restrictive initial schema determines what counts as "linguistic experience" and what knowledge arises on the basis of this experience. I also think, and have argued elsewhere, that the empiricist doctrines that have been prevalent in linguistics, philosophy, and psychology in recent years, if formulated in a fairly precise way, can be refuted by careful study of language. If philosophy is what philosophers do, then these conclusions are relevant to philosophy, both in its classical and modern varieties.

At this point, I would like to turn to some of the critical analysis of this point of view that has appeared in the recent philosophical literature, specifically, to the items referred to in note 2.

Goodman's treatment of these questions seems to me to suffer, first, from a historical misunderstanding; second, from a failure to formulate correctly the exact nature of the problem of acquisition of knowledge; and third, from a lack of familiarity with the work that has led to the conclusions that he criticizes, those that are outlined above.

His historical misunderstanding has to do with the issue between Locke and whoever it was that Locke thought he was criticizing in his discussion of innate ideas. Goodman believes that "Locke made . . . acutely clear" that the doctrine of innate ideas is "false or meaningless." I will not dwell on this matter, since it is a commonplace of historical scholarship that Locke's critique of the doctrine of innate ideas "assails it in its crudest form, in which it is countenanced by no eminent advocate."[4] Even Lord Herbert makes it clear that the common notions "remain latent" in the absence of appropriate stimulation, that they are the "principles without which we should have no experience at all" but that they will obviously not be constantly in consciousness, even to "normal men," and certainly not to those who are "headstrong, foolish, weak-minded and imprudent," to "madmen, drunkards, and infants," and so on. And as these ideas are elaborated by Descartes and others, it is repeatedly emphasized that while innate ideas and principles determine the nature of experience and the knowledge that can arise from it, they will ordinarily not be in consciousness. Since Locke's arguments fail to come to grips with the "dispositional" nature of innate structure that is insistently maintained by the leading proponents of rationalist doctrine, they also invariably miss the mark; it seems that he must have mistaken the actual views of Herbert, Descartes, the minor Cartesians, Cudworth, and others.

It is surprising that Goodman accuses those who "identify the innate ideas with capacities" of "sophistry." Goodman is free, if he wishes, to use the terms "idea" and "innate idea" in accordance with Locke's misunderstanding of rationalist doctrine, but hardly to accuse others of "sophistry" when they examine and develop this doctrine in the form in which it was actually presented. It is particularly surprising to hear Goodman speak of the necessity of applying the term "idea," in "its normal use." One would hardly expect Goodman to propose this sort of "ordinary language argument" against the use of a technical term. Furthermore, as Thomas Reid pointed out, if we use "idea" in the nontechnical way, then not only the position of Descartes, but also that of Locke and Hume reduces to absurdity – an observation that is correct, but that shows nothing more than the absurdity of insisting that a technical term must be understood in "the normal use" of the homonymous nontechnical term of ordinary discourse.

Let me turn, however, to the substantive problem of acquisition of knowledge, as Goodman formulates it in the specific case of language-acquisition. Quite properly, he distinguishes two cases: initial language, and second-language acquisition. But his analysis of the two cases leaves much to be desired.

Consider first the problem of second-language acquisition. In what I understand to be Goodman's view,[5] second-language acquisition poses no problem, since "once one language is available and can be used for giving explanation and instruction, the limitations [determined by an innate schematism] are transcended." This way of putting the matter misconstrues the situation in two basic respects. First, it is misleading to speak of the innate schematism that has been proposed as merely providing "limitations" for acquisition of language. Rather, what has been proposed is that this schematism ***makes possible*** the acquisition of a rich and highly specific system on the basis of limited data. To take one example, the problem is to explain how the data available to a language learner (first or second) suffices to establish that the phonological rules (the rules that assign phonetic representations to surface structures) apply cyclically, first to innermost phrases of the surface structure, then to larger phrases, etc., until the maximal domain of phonological processes – in simple cases, the full sentence – is reached. There is in fact good evidence that the rules do apply cyclically, but this evidence is not of a sort that can be used as the basis for induction from phonetic data to the principle of cyclic application, by any procedure of induction that has general validity. In particular, much of this evidence is derived from an analysis of percepts, that is, from investigation of the way in which someone who has already mastered the language interprets speech signals. It seems that this interpretation imposes a certain structure that is not indicated directly in the speech signal, for example, in the determination of stress contours.[6] Obviously the child cannot acquire the knowledge that phonological rules apply cyclically from data that are available to him only after he knows and makes use of this principle. This is an extreme example, but it nevertheless illustrates quite well the basic problem: to explain how a rich and highly specific grammar is developed on the basis of limited data that is consistent with a vast number of other conflicting grammars. An innate schematism is proposed, correctly or incorrectly, as an empirical hypothesis to explain the uniformity, specificity, and richness of detail and structure of the grammars that are, in fact, constructed and used by the person who has mastered the language. Therefore the word "limitation" in Goodman's formulation is quite inappropriate.

More serious, it must be recognized that one does not learn the grammatical structure of a second language through "explanation and instruction," beyond the most elementary rudiments, for the simple reason that no one has enough explicit knowledge about this structure to provide explanation and instruction. For example, consider the property of nominalization in English noted earlier, namely, that a certain class of nominal expressions correspond only to deep and not surface structures. The person who has learned English as a second language well enough to make the judgments illustrated by examples (1)–(10) has not acquired this knowledge through "explanation and instruction." Until quite recently, no one, to my knowledge, was aware of this phenomenon; the second-language learner, like the first-language learner, has somehow established the facts for himself, without explanation or instruction. Again, the example is quite typical. Only a trivial part of the knowledge that the second-language learner acquires is presented to him by direct instruction. Even the most cursory attention to the facts of second-language acquisition is sufficient to establish this. Hence, although second-language acquisition is, indeed, to be distinguished from first-language acquisition, the distinction is not of the sort that Goodman suggests. While it may be true that "once some language is available, acquisition of others is relatively easy," it nevertheless remains a very serious problem – not significantly different from the problem of explaining first-language acquisition – to account for this fact.

Consider now the more important matter of first-language acquisition, the problem to which the empirical hypotheses regarding innate schematism have been directed. Goodman argues that there is no problem in explaining first-language acquisition, because "acquisition of an initial language is acquisition of a secondary symbolic system": 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 some of the specific properties of grammar – say the distinction of deep and surface structure, the specific properties of grammatical transformations and phonological rules, the principles of rule ordering, and so on – were present in these already acquired prelinguistic "symbolic systems." But there is not the slightest reason to believe that this is so. Goodman's argument is based on a metaphorical use of the term "symbolic system," and collapses as soon as we try to give this term a precise meaning. If it were possible to show that "prelinguistic symbolic systems" share certain significant properties with natural language, we could then argue that these properties of natural language are somehow acquired by "analogy," though we would now face the problem of explaining how the "prelinguistic symbolic systems" developed these properties and how the analogies are established. But the issue is academic, since, for the moment, there is no reason to suppose the assumption to be true. Goodman's argument is a bit like a "demonstration" that there is no problem in accounting for the development of complex organs, because everyone knows that mitosis takes place. This seems to me to be obscurantism, which can be maintained only so long as one fails to come to grips with the actual facts.

There is, furthermore, a non sequitur in Goodman's discussion of first- and second-language acquisition. Recall that he explains the presumed ease of second-language acquisition on the grounds that it is possible to use the first language for explanation and instruction. He then goes on to argue that "acquisition of an initial language is acquisition of a secondary symbolic system," and is hence quite on a par with second-language acquisition. The primary symbolic systems he has in mind are "rudimentary prelinguistic symbolic systems in which gestures and sensory and perceptual occurrences of all sorts function as signs." But evidently, these systems, whatever they may be, cannot "be used for giving explanation and instruction" *in the way in which a first language can be used in second-language acquisition.* Consequently, 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, exactly as comparable experimental tests are not feasible in any other area of human psychology. But there is no reason for dismay at the impracticality of such direct tests as these. There are many other ways – those discussed earlier, and extensively in the literature – in which evidence can be obtained regarding the properties of grammars and in which hypotheses regarding the general properties of such grammars can be put to empirical test. Any such hypothesis immediately specifies, correctly or incorrectly, certain properties of "bad" languages. It therefore makes an empirical claim that can be falsified by finding counterinstances in some human language, or by showing that under the actual conditions of language acquisition, the properties in question do not appear in the system that is developed by the language learner. In linguistics, as in any other field, it is only in such indirect ways as these that one can hope to find evidence bearing on nontrivial hypotheses. Direct experimental tests of the sort that Goodman, for some reason, regards as necessary, are rarely feasible, a fact that may be unfortunate but that is nevertheless characteristic of most research.

Goodman's further claim, that not "a single general property of 'bad' languages has been formulated," is quite unfair. There are dozens of books and papers concerned with formulating properties of universal grammar and examining their empirical consequences, and each such property specifies "bad" languages, as just noted. One is free to argue that these attempts are misguided, inadequate, unconvincing, refuted by facts, etc., but not to deny blandly that they exist. I do not see how to avoid the conclusion that when Goodman speaks of "the unimpressive evidence adduced with respect to languages," he simply speaks out of ignorance, rather than from a considered analysis of the work that has been done in the field.

In discussing properties of "bad" languages, Goodman refers only to one case, namely, the case of the concocted language *Gruebleen*, which "differs from ordinary English only in that it contains the predicates 'grue' (for "examined before t and green or not so examined and blue") and 'bleen' (for "examined before t and blue or not so examined and green") in place of the predicates 'green' and 'blue.'" He argues that even in this case, one must be "painfully aware of the difficulties of answering" the question of what in general is "the difference between Gruebleen-like and English-like languages." I think that this is a rather marginal issue, since much more deep-seated properties of "English-like languages" have been formulated and investigated, but, since he brings up this example, it is well to point out that the difficulties to which he alludes are in large measure a consequence of the vagueness of the question he asks. Thus there is no difficulty in finding some property of Gruebleen that is not a property of "English-like languages," even a property of some generality. For example, consider the predicate "match" understood as in Goodman's *Structure of Appearance*, but applying now to objects rather than qualia. Thus two objects match "if and only if they are not noticeably different on direct comparison."[7] Gruebleen has the curious property that if an object A is examined before t and an object B is examined after t, and both are found to be grue (or both bleen), then we know that they will not match. But there is not t such that given two objects, one examined before t and one after t, and both found to be green (or blue), we can predict that they will not match. They may not match, but then they also may match, if both are green (or blue). In fact, it is undoubtedly a general property of natural languages that they are "English-like" rather than "Gruebleen-like," in this sense, in the domain of color terms. Thus there is no difficulty in establishing a fairly general distinction between Gruebleen-like and English-like languages, in this specific respect. Of course, this would not satisfy Goodman's requirements, for his special purposes, because one can construct other problems of the grue-bleen type that are not taken account of by this property. As long as Goodman's vague notions "English-like" and "Gruebleen-like" are left unspecified, there is of course no way to meet his demand that a *general* property be stated distinguishing the two kinds of language, and any specific distinction that is proposed will always give rise to new riddles of induction. This is an interesting comment about the limitations of inductive methods, but has no more relevance to the problem of specifying the characteristics of universal grammar than to any other enterprise of science, say, the problem of specifying the genetic conditions that determine that a human embryo will develop legs rather than wings, under a given range of conditions.

I am not, incidentally, proposing that the property just cited serves to explain why every language-learner (in fact, every mouse, chimpanzee, etc.) uses green rather than grue as the basis for generalization. No doubt this is a simple consequence of certain properties of the sensory system, a conclusion that is quite uninteresting from Goodman's point of view, but not, for that reason, incorrect.

Returning to the main point, it is interesting that at one stage of his argument Goodman remarks, quite correctly, that even if "for certain remarkable facts I have no alternative explanation," "that alone does not dictate acceptance of ... an intrinsically repugnant and incomprehensible theory." But now let us consider the theory of innate ideas that arouses Goodman's indignation, and ask whether it is "incomprehensible" and "repugnant."

Consider first the matter of comprehensibility. It does not seem to me incomprehensible that some aspect of the "final state" of an organism or automaton should also be an aspect of its "initial state," prior to any interaction with the environment, just as it is not incomprehensible that this aspect of the final state should have developed through internal processes, perhaps set in motion by organism–environment interaction of some sort. But consider the actual doctrines developed in the speculative psychology of rationalism, rather than Locke's caricature. Descartes, for example, argued that the idea of a triangle is innate in that "the idea of a true triangle ... can be more easily conceived by our mind than the more complex figure of the triangle drawn on paper," so that when a child first sees the more complex figure, he will "apprehend not it itself, but rather the authentic triangle." As Cudworth elaborates this view, "every irregular and imperfect triangle [is] as perfectly that which it is, as the most perfect triangle," but we interpret sensory images in terms of a notion of "regular figure" that has its source in the "rule, pattern and exemplar" generated by the mind as an "anticipation," just as we interpret all sensory data in terms of certain concepts of object and relations among objects, certain notions of cause and effect, Gestalt properties, functions in a "space" of possible human actions, and so on. Neither this view, nor its elaboration in modern psychology, is incomprehensible, though it may of course be misguided or incorrect. Similarly, there is no difficulty in comprehending the proposal that there are certain innate conditions on the form of grammar that determine what constitutes linguistic experience and what knowledge will arise on the basis of this experience. Again, one can easily design an automaton that will function in this manner, so that although the proposal may be wrong, it is not incomprehensible.

Whatever Goodman's attitudes might be to these formulations, it is interesting that he appears quite willing, at least in this paper, 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 terminology. His argument is directed not against the notion that "ideas are in the mind," but rather against the assumption that they are "in the mind" prior to experience, and surely if one assumption is comprehensible, then the other is as well (though neither, as noted, does justice to the classical rationalist view or to its modern variants). On the other hand, this approach to the problem of acquisition of knowledge will, no doubt, be "repugnant" to one who considers empiricist doctrine immune to doubt or challenge. But this is to treat empiricist doctrines as articles of religious faith. Surely it is not reasonable to be so bound to a tradition as to refuse to examine conflicting views about acquisition of knowledge on their merits.

Let me turn next to Hilary Putnam's contribution to the same symposium. Although his paper deals more directly with the points that are actually at issue, still it seems to me that his arguments are inconclusive, primarily, because of certain erroneous assumptions about the nature of the acquired grammars. Specifically, he enormously underestimates, and in part misdescribes, the richness of structure, the particular and detailed properties of grammatical form and organization that must be accounted for by a "language acquisition model," that are acquired by the normal speaker-hearer and that appear to be uniform among speakers and also across languages.

To begin with, Putnam assumes that at the level of sound structure, the only property that can be proposed in universal grammar is that a language has "a short list of phonemes." This uniformity among languages, he argues, requires no elaborate explanatory hypothesis. It can be explained simply in terms of "such parameters as memory span and memory capacity," and no "rank Behaviorists" would have denied that these are innate properties. In fact, however, very strong empirical hypotheses have been proposed regarding the choice of universal distinctive features, the form of phonological rules, the ordering and organization of these rules, the relation of syntactic structure to phonetic representation, none of which can conceivably be accounted for on grounds of memory limitations. Putnam bases his account largely on my "Explanatory Models in Linguistics" (see note 6), which examines in some detail the principle of cyclic application of phonological rules, a principle that, if correct, raises some rather serious problems. We must ask how the child acquires knowledge of this principle, a feat that is particularly remarkable since, as already noted, much of the evidence that leads the linguist to posit this principle is drawn from the study of precepts and is thus not even available to the child. Similar questions arise with respect to many other aspects of universal phonology. In any event, if the proposals that have been elaborated regarding sound structure are correct or near correct, then the similarities among languages at this level, and the richness of the knowledge acquired by the child, are indeed remarkable facts, and demand an explanation.

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 specific character 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 ... 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'"; that there is nothing surprising in the fact that languages contain rules of abbreviations. Hence, he concludes, the only innate conditions that must be postulated are those that apply to all reasonable "computing systems," and no Behaviorist should feel any surprise at this.

Each of the three conclusions, however, is vitiated by a false assumption. First, it is obvious that there are many different phrase-structure grammars consistent with the assumption that one of the categories is that of proper names. In fact, there is much dispute at the moment about the general properties of the underlying base system for natural languages; the dispute is not in the least resolved by the existence of proper names as a primitive category in many languages.[8]

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 have even an indirect relevance to this matter are those dealing with context-free phrase-structure grammars and their automata-theoretic interpretation. Context-free grammars are a reasonable model for the rules generating deep structures, when we exclude the lexical items and the distributional conditions they meet. But even apart from this fundamental discrepancy, the only existing results relate context-free grammars to a class of automata called "nondeterministic pushdown storage automata," and these have no particularly striking properties insofar as speed or complexity of computation are concerned, and are certainly not "natural" from this point of view. In terms of time and space conditions on computation, the somewhat similar but not formally related concept of real-time deterministic automaton would seem to be far more natural. In short, there are no results demonstrating that phrase-structure grammars are optimal in any computational sense (nor, certainly, are there any results dealing with the much more complex notion of base structure with a context-free phrase-structure grammar and a lexicon, with much richer properties, as components).

But there is no point in pursuing this matter, since what is at stake, in any event, is not the "simplicity" of phrase-structure grammars but rather of transformational grammars that contain a phrase-structure component, the latter playing a role in the generation of deep structures. And there is absolutely no mathematical concept of "ease of computation" or "simplicity of algorithm" that even suggests that such systems have some advantage over the various kinds of automata that have been investigated from this point of view. In fact, these systems have never really been considered in a strictly mathematical context, though there are interesting initial attempts to study some of their formal properties.[9] The source of the confusion is a misconception on Putnam's part as to the nature of grammatical transformations. These are not, as he supposes, rules that "abbreviate" sentences generated by phrase-structure rules. Rather, they are operations that form surface structures from underlying deep structures, which are generated, in part, by phrase-structure rules. Although there has been considerable evolution of theory since the notions of transformational generative grammar were first proposed, one assumption that has remained constant is that the phrase-structure rules generate only abstract structures, which are then mapped onto surface structures by grammatical transformations – the latter being structure-dependent operations of a peculiar sort that have never been studied outside of linguistics, in particular, not in any branch of mathematics with which I am familiar. To show that transformational grammars are the "simplest possible" one would have to demonstrate that an optimal computing system would take a string of symbols as input and determine its surface structure, the underlying deep structure, and the sequence of transformational operations that relate these two labeled bracketings. Nothing known about ease or simplicity of computation gives any reason to suppose that this is true; in fact, the question has never been raised. One can think of certain kinds of organization of memory that might be well adapted to transformational grammars, but this is a different matter entirely.[10] I would, naturally, assume that there is some more general basis in human mental structure for the fact (if it is a fact) that languages have transformational grammars; one of the primary scientific reasons for studying language is that this study may provide some insight into general properties of mind. Given those specific properties, we may then be able to show that transformational grammars are "natural." This would constitute real progress, since it would now enable us to raise the problem of innate conditions on acquisition of knowledge and belief in a more general framework. But it must be emphasized that, contrary to what Putnam asserts, there is no basis for assuming that "reasonable computing systems" will naturally be organized in the specific manner suggested by transformational grammar.

I believe that this disposes of Putnam's main argument, namely, that there is "nothing surprising," even to a Behaviorist, in the linguistic universals that are now being proposed and investigated. Let me then turn to his second argument, that even if there were surprising linguistic universals, they could be accounted for on a simpler hypothesis than that of an innate universal grammar, namely, the hypothesis of common origin of languages. This proposal misrepresents the problem at issue. As noted earlier, the empirical problem we face is to devise a hypothesis about initial structure rich enough to account for the fact that a specific grammar is acquired, under given conditions of access to data. To this problem, the matter of common origin of language is quite irrelevant. The grammar has to be discovered by the child on the basis of the data available to him, through the use of the innate capacities with which he is endowed. To be concrete, consider again the two examples discussed above: the association of nominal phrases to base structures and the cyclic application of phonological rules. The child masters these principles (if we are correct in our conclusions about grammar) on the basis of certain linguistic data; he knows nothing about the origin of language and could not make use of such information if he had it. Questions of common origin are relevant to the empirical problems we are discussing only in that the existing languages might not be a "fair sample" of the "possible languages," in which case we might be led mistakenly to propose too narrow a schema for universal grammar. This possibility must be kept in mind, of course, but it seems to me a rather remote consideration, given the problem that is actually at hand, namely, the problem of finding a schema rich enough to account for the development of the grammars that seem empirically justified. The discovery of such a schema may provide an explanation for the empirically determined universal properties of language. The existence of these properties, however, does not explain how a specific grammar is acquired by the child.

Putnam's discussion of the ease of language-learning seems to me beside the point. The question whether there is a critical period for language-learning is interesting,[11] but it has little relevance to the problem under discussion. Suppose that Putnam were correct in believing that "certainly ... 600 hours [of direct method instruction] will enable any adult to speak and read a foreign language with ease." We would then face the problem of explaining how, on the basis of this restricted data, the learner has succeeded in acquiring the specific and detailed knowledge that enables him to use the language with ease, and to produce and understand a range of structures of which the data presented to him constitute a minute sample.

Finally, consider the alternative approach that Putnam suggests to the problem of language acquisition. He argues that instead of postulating an innate schematism one should attempt to account for this achievement in terms of "general multi-purpose learning strategies." It is these that must be innate, not general conditions on the form of the knowledge that is acquired. Evidently, this is an empirical issue. It would be sheer dogmatism to assert of either of these proposals (or of some particular combination of them) that it *must* be correct. Putnam is convinced, on what grounds he does not say, that the innate basis for the acquisition of language must be identical with that for acquiring any other form of knowledge, that there is nothing "special" about the acquisition of language. A nondogmatic approach to this problem can be pursued, 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 such competence. If we discover that the same "learning strategies" are involved in a variety of cases, and that these suffice to account for the acquired competence, then we will have good reason to believe that Putnam's empirical hypothesis is correct. If, on the other hand, we discover that different innate systems (whether involving schemata or heuristics) have to be postulated, then we will have good reason to believe that an adequate theory of mind will incorporate separate "faculties," each with unique or partially unique properties. I cannot see how one 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."[12] Invoking an innate representation of universal grammar does solve the problem of learning (at least partially), in this case, if in fact it is true that this is the basis (or part of the basis) for language acquisition, as it well may be. If, on the other hand, there exist general learning strategies that account for the acquisition of grammatical knowledge, then postulation of an innate representation of 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 stages of investigation. At the moment, the only concrete proposal that is at all plausible, in my opinion, is the one sketched above. When some "general learning strategy" is suggested, we can look into the relative adequacy of these alternatives, on empirical grounds.

Henry Hiż's review article deals mainly with the distinction between competence and performance. One can attempt to explain technical concepts such as these in two different ways. At a presystematic level, one can try to indicate, necessarily in a loose and somewhat vague and only suggestive way, just what role the concept is intended to play in a more general framework, and why it seems to be a useful idea to try to develop. Discussion at this level is entirely legitimate, but there will generally be much room for misunderstanding. At a second level, one can develop the concept in as precise a way as the state of the field permits, with no consideration for motivation or general implications. At this level, the problem is to determine not what the concept in question is, but why there is any point in developing it.

At the presystematic level, I have tried to explain what I mean by "linguistic competence" in terms of models of use and acquisition of language, in the manner outlined earlier. At the systematic level, competence is expressed by a generative grammar that recursively enumerates structural descriptions of sentences, each with its phonetic, syntactic, and semantic aspects. It is hardly necessary to emphasize that any such grammar that we can actually present today is incomplete, not only because our knowledge of particular languages is deficient, but also because our understanding of phonetic and semantic representation and the kinds of structures and rules that mediate between them is limited and unsatisfactory in many respects.

Turning to Hiż's paper, there is, not surprisingly, a certain degree of misunderstanding between us at the presystematic level. Hiż suggests that my use of the notion "competence" "is to be understood as saying that introspection is a source of linguistic knowledge." I do agree that introspection is an excellent source of data for the study of language, but this conclusion does not follow from the decision to study linguistic competence. One might (irrationally, in my opinion) refuse to use such evidence, and still try to discover the generative grammar that represents "what is learned" and that plays a fundamental role in language use. This decision would be pointless, rather on a par with an astronomer's refusal, at one stage of the science, to use what he sees through a telescope as data, but the decision has nothing to do with the distinction between competence and performance. I have no doubt that it would be possible to devise operational and experimental procedures that could replace the reliance on introspection with little loss, but it seems to me that in the present state of the field, this would simply be a waste of time and energy. Obviously, any such procedure would first have to be tested against the introspective evidence. If one were to propose a test for, say, grammaticalness, that fails to make the distinctions noted earlier in the proper way, one would have little faith in the procedure as a test for grammaticalness. To me it seems that current research is not hampered significantly by lack of accurate data, but rather by our inability to explain in a satisfactory way data that are hardly in question. One who feels differently can support his point of view by demonstrating the gains in insight and understanding that can be achieved by refinements in techniques of data collection and analysis, say, by operational techniques for establishing grammaticalness, techniques that have been judged by the prior test of intuition and shown to be sufficiently sound so that one can rely on them in difficult or obscure cases. In any event, the whole matter has nothing to do with the decision to study linguistic competence.

Hiż regards it as "paradoxical" to assert, as I have, that linguistics "attempts to specify what the speaker actually knows, not what he may report about his knowledge." This he regards as "a peculiar sense of 'knowledge.'" To me it seems a rather ordinary sense, and a nonparadoxical usage. A person who knows English may give all sorts of incorrect reports about the knowledge that he actually possesses and makes use of constantly, without awareness. As noted earlier, when we study competence – the speaker-hearer's knowledge of his language – we may make use of his reports and his behavior as evidence, but we must be careful not to confuse "evidence" with the abstract constructs that we develop on the basis of evidence and try to justify in terms of evidence. Thus I would definitely reject three of the five conditions that, Hiż suggests, rules must satisfy if they are to constitute an account of competence in my sense, namely, that the native speaker feels that the sentences generated by the rules are in his language, that they have the assigned structures, and that what the speaker feels is true. Since performance – in particular, judgments about sentences – obviously involves many factors apart from competence, one cannot accept as an absolute principle that the speaker's judgments will give an accurate account of his knowledge. I am surprised that Hiż should offer this interpretation of my views immediately after having quoted my statement that the speaker's reports about his competence may be in error.

At least for the purposes of discussion, Hiż is willing to accept the view that a generative grammar, a system of rules assigning structures to sentences, can serve to characterize competence. He then points out, correctly, that the linguist is guided in his choice of a grammar by certain "general principles about language as such," and that this general theory – universal grammar – will have explanatory value if it selects particular grammars correctly. He then attributes to me, incorrectly, the view that universal grammar is to be identified with "a theory of language acquisition." My view, rather, is that universal grammar is one element of such a theory, much as competence is one element of a theory of performance. There are surely many other factors involved in language acquisition beyond the schematism and weighting function that – if my suggestion is correct – play a part in determining the nature of the acquired competence. This misinterpretation of my proposal regarding the relation of universal grammar to language acquisition parallels the misinterpretation of my proposal regarding the relation of competence to performance; in both cases, what is omitted is the reference to other factors that must be involved. In the case of language acquisition, furthermore, it must be emphasized that the model I am suggesting can at best only be regarded as a first approximation to a theory of learning, since it is an instantaneous model and does not try to capture the interplay between tentative hypotheses that the child may construct, new data interpreted in terms of these hypotheses, new hypotheses based on these interpretations, and so on, until some relatively fixed system of competence is established. I think that an instantaneous model is a reasonable first approximation, but this, as any other aspect of research strategy, must ultimately be evaluated in terms of its success in providing explanations and insight.

Hiż regards the reference to classical formulations of problems of language and mind as "confusing and misleading historical baggage." I disagree with this judgment, but have nothing to add here beyond what I have written elsewhere.[13] My feeling is that the contributions of rationalist psychology and linguistics are interesting in themselves, and are quite relevant to present concerns, more so, in fact, than much of the work of the past century. One who finds these forays into intellectual history "confusing and misleading" can perfectly well disregard them. I see no issue here.

Before leaving this matter, I should mention that Hiż is inaccurate in stating that Herbert of Cherbury restricted himself to "religious knowledge." Nor can Thomas Reid be described as one of those concerned to develop a doctrine of innate universals. Furthermore, it is surely misleading to say that I "call upon" Descartes and others "to support" my "stand on innate universals." Their advocacy of a similar position does not constitute "support." Rather, I am suggesting that their contributions have been inadequately appreciated, and that we can still learn a good deal from a careful study of them.

Hiż objects to the fact that my proposals concerning universal grammar are based on detailed examination of a few languages rather than "examination of many cases." I certainly agree that one should study as many languages as possible. Still, a *caveat* should be entered. It would be quite easy to present enormous masses of data from varied languages that are compatible with all conceptions of universal grammar that have so far been formulated. There is no point in doing so. If one is concerned with the principles of universal grammar, he will try to discover those properties of particular grammars that bear on these principles, putting aside large amounts of material that, so far as he can determine, do not. It is only through intensive studies of particular languages that one can hope to find crucial evidence for the study of universal grammar. One study such as that of Matthews on Hidatsa (see note 8) is worth one thousand superficial studies of varied languages from this point of view. If someone feels that the base of data is too narrow, what he should do is show that some of the material omitted refutes the principles that have been formulated. Otherwise, his criticism has no more force than a criticism of modern genetics for basing its theoretical formulations on the detailed investigation of only a few organisms.

Hiż also argues that the principles of universal grammar, even if true, may indicate only "the common historical origin of languages." I have already pointed out why this hypothesis is without explanatory force.

Hiż maintains that decisions about particular parts of grammar (by the linguist) are "determined not by a general theory but by internal usefulness within the particular grammar," and objects that I do not make this clear. Since I have no idea what is meant by "internal usefulness," I have nothing to say about this point. The issue is confused by his misinterpretation of my use of the notion "simplicity." When I speak of "simplicity of grammar," I am referring to a "weighting function," empirically determined, that selects a grammar of the form permitted by the universal schematism over others that are also of the proper form and are compatible with the empirical data. I am not using the term "simplicity" to refer to that poorly understood property of theories that leads the scientist to select one rather than another. The evaluation measure that defines "simplicity of grammars" is part of linguistic theory. We must try to discover this measure on empirical grounds, by considering the actual relations between input data and acquired grammars. Thus the notion "simplicity of grammar" plays a role analogous to that of a physical constant; we must establish it on empirical grounds, and there is no a priori insight on which we can rely. The problems of defining "simplicity of theories" in a general context of epistemology and philosophy of science are entirely irrelevant to the issue of determining, on empirical grounds, the properties of grammars that lead to the selection of one rather than another in language acquisition. This point has been emphasized repeatedly. See, for example, *Aspects*, chapter 1, section 7.

One final comment. Hiż suggests that "it should be easier to explain why we assign such-and-such a structure to a sentence by pointing out how this sentence changes the readings of neighboring sentences than by referring to innate universal ideas and mental reality." Here he is confusing two entirely different kinds of explanation. If I want to explain why, yesterday afternoon at three o'clock, John Smith understood "the shooting of the hunters" as referring to the act of shooting the hunters, rather than the hunters' act of shooting, I will of course bring into consideration the situational context (not limiting myself to the "readings of neighboring sentences"). If I am interested in explaining why this phrase is suspectible to these two interpretations, but the phrase "the growth of corn" is susceptible to only one (namely, the corn's growing and not the act or process of growing corn), then I will appeal first to the particular grammar of English, and more deeply, to the linguistic universals that led to the construction of this grammar by a child exposed to certain data. Since entirely different things are being explained, it is senseless to claim that one manner of explanation is "easier" than the other.

Harman's critique is also concerned with the matter of competence and performance. He begins by ascribing to me a view that I have never held, and have explicitly rejected on numerous occasions, namely, that "competence [is] the knowledge that the language is described by the rules of the grammar," and that a grammar describes this "competence." Obviously, it is absurd to suppose that the speaker of the language knows the rules in the sense of being able to state them. Having attributed to me this absurd view, Harman goes on to struggle with all sorts of purported confusions and difficulties of interpretation. But he cites nothing that could possibly be regarded as a basis for attributing to me this view, though he does quote remarks in which I explicitly reject it. Therefore, I will not discuss this part of his argument at all.

In Harman's framework, there are two kinds of knowledge: knowing that and knowing how. Obviously knowledge of a language is not a matter of "knowing that." Therefore, for him, it must be a matter of "knowing how." A typical speaker "knows how to understand other speakers"; his competence is his ability "to speak and understand the language described by [the] grammar" that describes the language. I do not know what Harman means by the locution "knows how to understand," but clearly he is using the term "competence" in a different way from what I proposed in the work he is reviewing. In my sense of "competence," the ability to speak and understand the language involves not only "competence" (that is, mastery of the generative grammar of the language, tacit knowledge of the language), but also many other factors. In my usage, the grammar is a formal representation of what I have called "competence." I have no objection to Harman's using the term in a different way, but when he insists on supposing that his usage is mine, naturally, only confusion will result. Again, I see no point in tracing in detail the various difficulties into which this misinterpretation leads him.

According to Harman, the "competence to speak and understand the language" is a skill, analogous to the skill of a bicycle rider. Given his insistence that knowledge of language is a matter of "knowing how" (since it is obviously not "knowing that"), this is not an unexpected conclusion. But he suggests no respect in which ability to use a language (let alone the competence, in my sense, that constitutes an element of this ability) is like the ability to ride a bicycle, nor do I see any. The proper conclusion, then, would be that there is no reason to suppose that knowledge of language can be characterized in terms of "knowing how." I therefore see no point in the analogy that he suggests. Knowledge of language is not a skill, a set of habits, or anything of the sort. I see nothing surprising in the conclusion that knowledge of language cannot be discussed in any useful or informative way in this impoverished framework. In general, it does not seem to me true that the concepts "knowing how" and "knowing that" constitute exhaustive categories for the analysis of knowledge. Nor is it surprising that Harman finds it difficult to understand my remarks, or those of anyone else who is concerned with knowledge of language, given that he insists on restricting himself to this framework.

Harman tries to show that there is a fundamental incoherence in my proposal that in acquiring or using knowledge of a language (in developing "an internal representation of a generative system" or making use of it in speaking or understanding speech), the child makes use of an innate schematism that restricts the choice of grammars (in the case of acquisition) or an internalized grammar (in the case of language use). His argument seems to me unclear. As I understand it, it seems to proceed as follows. He argues that this internalized system must be presented in "another more basic language," which the child must come to understand before he can make use of this schematism to learn this language, or before he can make use of the grammar to understand speech. But this, he argues, leads to a vicious circle or an infinite regress. Thus if we were to say that the child knows the "more basic language" directly, without learning, then why not say also that he knows "directly the language he speaks," without learning; a vicious circle. Or, if we say that he must learn the more basic language, then this raises the question how the more basic language is learned, and leads to an infinite regress. This argument is totally invalid. Consider the case of acquisition of language. Even if we assume that the innate schematism must be represented in an "innate language," neither conclusion follows. The child must know this "innate language," in Harman's terms, but it does not follow that he must "speak and understand it" (whatever this might mean) or that he must learn it. All that we need assume is that he can make use of this schematism when he approaches the task of language-learning. So much for the infinite regress. As to the vicious circle, there is a very simple reason why we cannot assume that the child knows the language he speaks directly, without learning, namely, that the assumption is false. We cannot claim that every child is born with a perfect knowledge of English. On the other hand, there is no reason why we should not suppose that the child is born with a perfect knowledge of universal grammar, that is, with a fixed schematism that he uses, in the ways described earlier, in acquiring language. This assumption may be false, but it is quite intelligible. If one insists on describing this knowledge as "direct knowledge of a more basic language," I see no reason to object, so long as we are clear about what we mean, but would merely point out that there is no reason at all to doubt that the child has this direct knowledge. Hence there is no vicious circle, and no infinite regress. Similarly, if we consider the case of language use, there is neither incoherence nor implausibility. There is surely no infinite regress and no vicious circle in the assumption that in language use (speaking or understanding) the user employs an internally represented grammar. We can easily construct a model (say, a computer program) that functions in this way. I therefore fail to see any basis for Harman's belief that there is an infinite regress or vicious circle inherent in, or even suggested by this formulation.

In the second part of his paper, Harman turns to my argument that current work in linguistics supports a view of language and mind that has a distinctly rationalist flavor, and is in conflict with the empiricist views that have dominated the study of language and mind in recent years. He asserts that to infer a grammar from data, a model of language-learning must already have detailed information about the theory of performance. This is an interesting proposal, and it deserves to be developed. But I cannot go along with his rather dogmatic claim, hardly argued in the paper, that this approach must necessarily be correct, and that any other approach must fail to provide any insight into the problem of acquisition of knowledge. I think that the work of the past few years on universal grammar does, in fact, suggest and in part support an interesting, rather classical approach to the problem of how knowledge is acquired. In the absence of any argument as to why this approach must fail to be illuminating, I see no reason not to continue with the investigation of how principles of universal grammar might select a particular grammar on the basis of the data available.

Let us turn now to the issue of rationalist and empiricist approaches to problems of language and mind. As Harman points out, if we describe an innate schematism biased toward (or restricted to) a specific form of grammar as part of the "principles of induction used," and define "resourceful empiricism" as a doctrine that makes use of such "principles of induction" as this, then surely "resourceful empiricism" cannot be refuted, "no matter what the facts about language [or anything else] turned out to be." Of course, this new doctrine of "resourceful empiricism" would now incorporate "principles of induction" that are, so it seems, quite specific to the task of language acquisition and of no general validity.

The concept "resourceful empiricism" so defined seems to me of little interest. The issue that concerns me is whether there are "ideas and principles of various kinds that determine the form of the acquired knowledge in what may be a rather restricted and highly organized way," or alternatively, whether "the structure of the acquisition device is limited to certain elementary peripheral processing mechanisms ... and certain analytical data-processing mechanisms or inductive principles." (*Aspects*, pp. 47f.) I have argued that "it is historically accurate as well as heuristically valuable to distinguish these two very different approaches to the problem of acquisition of knowledge," even though they of course "cannot always be sharply distinguished" in the work of a particular person. (Ibid., p. 52.) In particular, I have tried to show that it is possible to formulate these approaches so that the former incorporates the leading ideas of classical rationalism as well as the modern variant I have been describing, and that the latter includes classical empiricist doctrine as well as the theories of acquisition of knowledge (or belief, or habit) developed in a wide range of modern work (Quine's notions of quality space and formation of knowledge by association and conditioning; Hull's approach in terms of primitive unconditioned reflexes, conditioning, and habit structures; taxonomic linguistics, with its analytic procedures of segmentation and classification and its conception of language as a "habit system," and so on).[14] Needless to say, there is no necessity to view the various attempts to study language acquisition within this framework; I can only say that I think it is both useful and accurate. These alternatives can be made fairly precise and investigated in terms of their empirical consequences. Harman's proposal to define "resourceful empiricism" in such a way as to include both approaches, and to be, as he notes, immune to any factual discovery, is merely a pointless terminological suggestion and cannot obscure the difference between the approaches mentioned or the importance of pursuing and evaluating them.[15]

To summarize, I doubt that linguistics can provide "a new technique" for analytic philosophy that will be of much significance, at least in its present state of development. Nevertheless, it seems to me that the study of language can clarify and in part substantiate certain conclusions about human knowledge that relate directly to classical issues in the philosophy of mind. It is in this domain, I suspect, that one can look forward to a really fruitful collaboration between linguistics and philosophy in coming years.

Notes

1 See, for example, my contribution to the symposium on innate ideas published in *Synthese*, XVII, 1 (March 1967), 2–11, and the references cited there on p. 11.

2 Specifically, the contributions by Nelson Goodman and Hilary Putnam to the symposium cited in note 1, pp. 12–28, and the review articles by Henry Hiż and Gilbert Harman in the issue of the *Journal of Philosophy* devoted to "Some Recent Issues in Linguistics," LXIV, 2 (February 2, 1967), 67–87. The latter two are largely devoted to critical analysis of chapter 1 of my *Aspects of the Theory of Syntax* (MIT Press, 1965).

3 Since the language has no objective existence apart from its mental representation, we need not distinguish between "system of beliefs" and "knowledge," in this case.

4 A. C. Fraser (ed.), in his edition of Locke's *Essay Concerning Human Understanding* (1894, republished by Dover, 1959), p. 38 of the Dover edition.

5 Cf. p. 24. Given the dialogue form of his article, it is difficult to be certain that one is not misrepresenting his position. However, I see no other way to interpret these remarks.

6 For some discussion, see my paper "Explanatory Models in Linguistics," in E. Nagel, P. Suppes, and A. Tarski (eds), *Logic, Methodology, and Philosophy of Science* (Stanford University Press, 1962). For some recent and much more extensive discussion, see N. Chomsky and M. Halle, *Sound Patterns of English* (Harper and Row, New York, 1968) and the references cited there, and my paper "Some General Properties of Phonological Rules," *Language*, XLIII (March 1967), 102–28.

7 N. Goodman, *Structure of Appearance* (2nd edition, Bobbs-Merrill), p. 272. The distinction between Gruebleen and English that I am now discussing is not to be confused with a pseudodistinction, correctly rejected by J. Ullian on the basis of a different usage of the notion "match." See *Phil. Rev.*, July 1961.

8 Not, incidentally, in all. Although this is hardly important here, it seems that many languages do not have proper names as a primitive category, but rather form proper names by recursive processes of an elaborate sort. See, for example, G. H. Matthews, *Hidatsa Syntax* (Mouton, The Hague, 1965), pp. 191ff.

9 See, for example, S. Peters and R. Ritchie, "On the Generative Capacity of Transformational Grammars," *Information and Control* (forthcoming), and J. P. Kimball, "Predicates Definable over Transformational Derivations by Intersection with Regular Languages," *Information and Control*, II (1967), pp. 177–95.

10 For some speculations on this matter, see G. A. Miller and N. Chomsky, "Finitary Models of Language Users," part II, in R. D. Luce, R. Bush, and E. Galanter (eds), *Handbook of Mathematical Psychology*, vol. II (Wiley, 1963).

11 See E. H. Lenneberg, *Biological Foundations of Language* (Wiley, 1967), for evidence bearing on this issue.

12 Or for his assumption that the "weighting functions" proposed in universal grammar constitute the "sort of fact . . . [that] . . . learning theory tries to account for; *not* the explanation being sought." No one would say that the genetic basis for the development of arms rather than wings in a human embryo is "the kind of fact that learning theory tries to account for," rather than the basis for explanation of other facts about human behavior. The question whether the weighting function is learned, or whether it is the basis for learning, is an empirical one. There is not the slightest reason to assume, a priori, that it is to be accounted for by learning rather than genetic endowment, or some combination of the two.

There are other minor points in Putnam's discussion that call for some comment. For example, he asserts that since certain ambiguities "require coaching to detect," it follows that "the claim that grammar 'explains the ability to recognize ambiguities' . . . lacks the impressiveness that Chomsky believes it to have." But he misconstrues the claim, which relates to competence, not performance. What the grammar explains is why "the shooting of the hunters" (the example he cites) can be understood with hunters as subject or object but that in "the growth of corn" we can understand "corn" only as subject (the explanation, in this case, turns on the relation of nominalizations to deep structures, noted earlier). The matter of coaching is beside the point. What is at issue is the inherent sound–meaning correlation that is involved in performance, but only as one of many factors. Putnam also misstates the argument for assuming the active–passive relation to be transformational. It is not merely that the speaker knows them to be related. Obviously that would be absurd; the speaker also knows that "John will leave tomorrow" and "John will leave three days after the day before yesterday" are related, but this does not imply that there is a transformational relation between the two. Syntactic arguments are given in many places in the literature. See, for example, my *Syntactic Structures* (Mouton, 1957); *Aspects of the Theory of Syntax*.

13 In my *Current Issues in Linguistic Theory* (Mouton, 1964), §1; *Aspects of the Theory of Syntax*, chapter 1, §8; *Cartesian Linguistics* (Harper and Row, 1966).

14 Harman observes correctly that I ignore the "enormous philosophical literature on induction," and limit myself solely to an investigation of the procedures of taxonomic linguistics as "the only proposals that are explicit enough to support serious study." He does not, however, show how anything in the literature on induction bears on the problems I am considering. The reason is that there is nothing. The literature on induction is quite interesting, but it happens to deal with entirely different questions. It does not even hint at procedures of analysis or acquisition of belief or confirmation that would overcome the problems that I have been discussing. There is, for example, nothing in the literature on induction that gives any insight into how the principles cited above as examples (the cycle of phonological rules or the rule of nominalization) might be reached "by induction" from the data available. But it is such questions as these that must be faced in the study of language acquisition.

15 Two minor points in this connection. Harman sees only a "tenuous historical connection" between procedures of segmentation and classification and phrase-structure grammar. The connection is actually much closer. Zelig Harris, in his *Methods of Structural Linguistics*, tried to show how a systematic use of such procedures, amplified by a simple inductive step, would lead to a set of rules that might be regarded as generating an infinite set of sentences. A set of Harris' "morpheme to utterance" formulas, though not quite the same as phrase-structure grammar, is quite similar. The concept of "phrase-structure grammar" was explicitly designed to express the richest system that could reasonably be expected to result from the application of Harris-type procedures to a corpus. Harris, and other methodologists of the 1940s, were developing an approach to linguistic analysis that one can trace at least to Saussure.

Secondly, Harman is quite correct in pointing out that in my reference to "the only [empiricist] proposals that are explicit enough to support serious study," I omitted mention of Harris' and Hiż's method of studying co-occurrence relationships. He feels that this method is "similar in spirit to the taxonomic procedures." I don't see the point in arguing this, one way or another. In any event, I know of no reason to suppose that such procedures can lead to or can even provide evidence for or against the postulation of a generative grammar.