THE JOURNAL OF PHILOSOPHY

VOLUME LXIX, NO. 22, DECEMBER 7, 1972

GRAMMAR, PSYCHOLOGY, AND INDETERMINACY

Significance is the trait with respect to which the subject matter of linguistics is studied by the grammarian.
Pending a satisfactory explanation of the notion of meaning, linguists in semantic fields are in the position of not knowing what they are talking about.

—W. V. Quine

ACCORDING to Quine, the linguist qua grammarian does not know what he is talking about. The goal of this essay is to tell him. My aim is to provide an account of what the grammarian is saying of an expression when he says it is grammatical, or a noun phrase, or ambiguous, or the subject of a certain sentence. More generally, I want to give an account of the nature of a generative grammatical theory of a language—of the data for such a theory, the relation between the theory and the data, and the relation between the theory and a speaker of the language.

I

Prominent among a linguist's pronouncements are attributions of grammaticality. What are we saying about a sentence when we say it is grammatical? One strategy for answering this question is to attend to the work of the grammarian. To be grammatical, a sentence must have those characteristics which the grammarian seeks in deciding whether a sentence is grammatical. So a reconstruction of the grammarian's work is a likely path to an explication of 'grammatical'. This is the strategy adopted by Quine,[1] and it will be of value to study his remarks in some detail. On Quine's account, *significance* rather than *grammaticality* "is the trait with respect to which the subject matter of linguistics is studied by the gram-

[1] *From a Logical Point of View*, 2d ed., revised (New York: Harper & Row, 1963), essay III.

marian" (48). If the two are different, there is some inclination to take the grammarian at his name. So let us see what can be learned by taking Quine's proposal as an explication of *grammaticality*.

The problem for the grammarian may be posed as the segregating of a class K of sequences that we will call *grammatical*. On Quine's view, he attends to four nested classes of sequences, H, I, J, and K.

> H is the class of observed sequences, excluding any which are ruled inappropriate in the sense of being non-linguistic or belonging to alien dialects. I is the class of all such observed sequences and all that ever will happen to be professionally observed, excluding again those which are ruled inappropriate. J is the class of all sequences ever occurring, now or in the past or future, within or without professional observation—excluding, again, only those which are ruled inappropriate. K, finally, is the infinite class of all those sequences, with the exclusion of the inappropriate ones as usual, which *could* be uttered without bizarreness reactions. K is the class which the grammarian wants to approximate in his formal reconstruction (53).

The linguist's data are H, and he checks his predictions against I minus H hoping that this will be a representative sample of J. It is when we come to K that philosophical eyebrows are raised; for what is the force of the 'could' which extends the class beyond J, commonly infinitely beyond? Quine's answer is that, besides H and future checks against I, the 'could' is the reflection of the scientist's appeal to simplicity. "Our basis for saying what 'could' be generally consists . . . in what *is* plus simplicity of the laws whereby we describe and extrapolate what is" (54).

Quine's proposal shares with other operational definitions the virtue of objectivity. Yet his solution is beset with problems. For Quine's procedure just does not pick out anything like the class we would pre-systematically hold to be grammatical—and this because his account fails to portray what the grammarian *actually does*. To see this, consider the case of a Quinean linguist ignorant of English setting out to segregate grammatical English sequences. He starts with H, the class of sequences he observes. But H, in addition to samples of what we would pre-systematically hold to be grammatical sequences, contains all manner of false starts, "lost thoughts," peculiar pauses ('aahhhh'!) and, unless he is uncommonly fortunate, a liberal sprinkling of blatantly incoherent speech. Yet Quine, if we take him literally, would have H included as a subset of K. What the resulting projection might be is hard to imagine. But K, so constructed, would not be the class of grammatical sequences in English.

It might be thought that, appealing to simplicity, the linguist could toss out an occasional member of *H*, much as he excludes from *H* what he takes to be nonlinguistic noise or intrusion from another tongue. But an hour spent attending carefully to unreflective speech will dispel this notion. There is simply too much to exclude.[2]

Quine succeeds in muddying the waters a bit by sprinkling the restriction that the sentences to be studied are those which could be uttered "without bizarreness reactions." It is not clear whether he takes such sentences to be excluded from *H* and *I* by virtue of their being observed *in situ* or whether he would have *H* and *I* further filtered. But it seems clear that, in either case, either this move is inadequate or it begs the question. If by 'bizarreness' Quine means *bizarreness,* then the exclusion will hardly accomplish his purpose. For many sorts of sequences that we would want to exclude from *K* (those with 'aahhh's' interspersed, for example, or those which change subject mid-sentence) are uttered all the time without bizarreness reactions. And many sentences we would want to include in *K* would surely evoke the strongest of bizarreness reactions. Indeed, though *K* will be infinite, only members of a finite subset could be uttered without evoking a bizarreness reaction. Sentences that take more than six months to utter are bizarre. If, however, the reaction Quine has in mind is the reaction (whatever it may be) characteristically displayed when an ungrammatical sequence is uttered, then, until he has provided some account of how this reaction is to be recognized, he has begged the question.[3]

II

Taking Quine's proposal as an explication of grammaticality has led to an impasse. In seeking our way around it we might do well to return to Quine's original insight and attend more closely to what the grammarian actually does. From the first, the generative grammarian has relied heavily on the fact that, with a modicum of instruction, speakers can be brought to make all manner of judgments about their language. In particular, they can be brought to make firm judgments on the oddness or acceptability of indefinitely many sequences. Provided with a few examples, speakers

[2] Much the same point is made by Jerrold Katz and Jerry Fodor in "What's Wrong with the Philosophy of Language?," *Inquiry,* v (1962): 197–237.

[3] Significance is likely a more inclusive notion than grammaticality, more liberal in the constructions it will allow and tolerating a richer sprinkling of 'aahhh's, 'I mean's, and 'you know's. Thus perhaps Quine's proposal does rather better when taken as advertised. But whatever its interest, significance as characterized by Quine is not the property studied by grammarians of a generative bent.

can go on to judge new sequences in point of grammaticality, and do so with considerable consistency for large numbers of cases. This suggests that we might try to remedy the difficulties with Quine's proposal by substituting *intuitive judgments* for observed utterances. On the revised account, H would be the class of those sequences which to date have been considered and judged to be grammatical. I would be the class of sequences ever reflected upon and judged clearly grammatical. And K is the infinite class projected along simplest lines from H and checked against I.

This modified account nicely circumvents the major shortcoming we found in Quine's proposal. Read literally, Quine's method did not pick out the class of sequences we would pre-systematically call grammatical. The class H on which his projection was based was already tainted with ungrammatical sequences. Our modified version avoids this difficulty by basing its projection on sequences intuitively taken to be grammatical. The projected class K can still miss the mark, failing to be compatible with I minus H. But this potential failure is the normal inductive one.[4]

We can now make a plausible first pass at depicting the grammarian's work. He proceeds by eliciting intuitive judgments about which sequences are in the informant's language and which are not. He then projects these clear cases along simplest lines, checking his projected class against speakers' intuitions. Thus the task of the generative grammarian may be viewed as that of constructing a system of rules and a definition of 'generate' that define a terminal language containing phonetic representations for all the sequences judged by speakers to be clearly acceptable and containing no sequence judged to be clearly unacceptable. The sequences about which speakers have no firm or consistent intuitions can be relegated to the class of "don't cares" and decided by the simplest grammar that handles the clear cases.

Yet as it stands the account still will not do. One fault is its myopic concentration on intuitions. Speakers' judgments about acceptability are the most important data for the grammarian. But they are not his only data, nor are they immune from being corrected or ignored. The attentive grammarian will attend to many aspects of his subjects' behavior in addition to their response to questions about sentences' acceptability. And a proper explication of the grammarian's job must provide some account of the role these additional data play.

[4] Note that Quine's "bizarreness reactions" could be taken as negative judgments when the subject is queried about a sequence's acceptability. If this is Quine's intention, his proposal and the present account converge.

Perhaps the most important sort of evidence for the grammarian besides intuitions of acceptability is the actual unreflective speech of his subjects. An informant's protest that a given sequence is unacceptable may be ignored if he is caught in the act, regularly uttering unpremeditatedly what, on meditation, he alleges he doesn't say. In addition to actual speech, there is a host of further clues for the grammarian. Stress patterns, facts about how sentences are heard and data on short-term verbal recall are among them.[5] Others might be mentioned. To what use does the grammarian put this further evidence? Principally, I suggest, to shore up the evidence provided by speakers' intuitive judgments or to justify his neglect of them. A sentence whose acceptability to speakers is in some doubt will, with good conscience, be generated by a grammar if it ranks high in the other tests. And, on the other side, a sentence that has the blessings of speakers may be rejected—not generated by the grammar—if it fails to display the other characteristics of grammatical sequences.

We now have one justification the grammarian may use for rejecting speakers' intuitions. There is another. And consideration of it will lead to a fundamental revision of our account of grammaticality. Intuitive oddness may be explained by many factors. Some sentences seem odd because they are pragmatically odd, describing a situation that is bizarre. Others, perhaps, may be rejected as obscene or taboo. Most importantly, sentences may seem odd because they are simply too long and complicated. If the grammarian suspects that any of these factors explain speakers' rejection of a sentence, he may classify it as grammatical *even though it lacks all the characteristics in the cluster associated with grammaticality.*

Note that at this juncture two notions we have been conflating part company. Thus far I have been interchanging 'acceptability' and 'grammaticality' with studied equivocation. Intuitions of acceptability and the cluster of further characteristics usually accompanying sentences judged acceptable have been taken as (more or less) necessary and sufficient conditions for grammaticality. But the picture changes when a sentence may be classed as grammatical in spite of failing each relevant test. The motivation for separating acceptability and grammaticality is *broad theoretic simplicity.* It is simpler to generate an infinite class including the acceptable sentences than it is to draw a boundary around just those sentences which rank high in the several tests for acceptability. But in thus

[5] Cf. George A. Miller and Stephen Isard, "Some Perceptual Consequences of Linguistic Rules," *Journal of Verbal Learning and Verbal Behavior,* II (1963): 217–228.

choosing the simpler task we must assume that some further theory or theories will account for those grammatical sentences which are unacceptable. And we must also assume that the new theory combined with a grammatical theory will together be simpler than any theory attempting directly to generate all and only the acceptable sequences. In short, we are venturing that the best theory to account for *all* the data will include a grammar of infinite generative capacity. This is hardly a step to be taken lightly. For in allowing his grammar to generate an infinite number of sentences, the grammarian is countenancing as grammatical an infinite number of sentences that fail each test of acceptability. It might be thought that such prodigality could be avoided by simply cutting off the class of sentences generated by a grammar at an appropriately high point. But this is not the case. For there is no natural point to draw the line—no point at which the addition of another conjunct or another clause regularly changes a clearly acceptable sentence into a clearly unacceptable one. Nor would it do to pick an *arbitrary* high cut-off point. This would leave the grammarian as before with generated sentences that are unacceptable. And any account of *why* these sentences were unacceptable would likely also account for the sequences beyond the arbitrary cut-off point.

By now it is evident that grammaticality is best viewed as a *theoretical* notion. Like other theoretical notions, it is related to relevant data in several and complex ways. Simple grammatical sentences generally have several or all of the cluster of characteristics typical of acceptable sequences. More complex grammatical sentences may share none of these characteristics. They are grammatical in virtue of being generated by the grammar that most simply generates all the clearly acceptable sentences and holds the best promise of fitting into a simple total theory of acceptability.

There is, thus, a conjecture built into a proposed grammar—the conjecture that this generative system will fit comfortably into a total theory that accounts for all the data. In this respect a grammar is similar to the theory of ideal gases. The ideal-gas laws do a good job at predicting the behavior of light gases at high temperatures and low pressures. In less favorable cases, the laws predict poorly. They were acceptable in the hope, later fulfilled, that further laws could be found to explain the difference between the behavior of real gases and the predicted behavior of ideal ones. The adoption of a given grammar or form of grammar might be viewed as setting up a "paradigm"[6] or framework for future investigation. The grammar

[6] In a sense that may be intended by T. S. Kuhn, *The Structure of Scientific Revolutions* (Chicago: University Press, 1962).

serves to divide those phenomena still needing explanation (viz., unacceptable grammatical sequences) from those already adequately handled.

In our portrait of the grammarian's job, the emphasis has shifted from the concept of grammaticality to the notion of a correct grammar. A sequence is grammatical if and only if it is generated by a correct grammar for the language in question. And a grammar is correct only if it excels in the virtues lately adumbrated. But there are higher virtues to which a grammar may aspire, and more data to be reckoned with. So far we have taken into account data about speakers' intuitions of acceptability and data about a cluster of further characteristics common among acceptable sequences. But we have hardly exhausted the speaker's intuitions about matters linguistic. There is a host of other properties of sentences and their parts about which speakers have firm intuitions. With a bit of training speakers can judge pairs of sentences to be related as active and passive, or as affirmative and negative. They can pick out parts of speech, detect subjects and objects, and spot syntactic ambiguities. The list of these grammatical intuitions could easily be extended. A grammatical theory will not only try to specify which sequences are acceptable; it will also try to specify the grammatical properties and relations of sentences as intuited by speakers. As in the case of intuitions of acceptability, the grammatical theory will be expected to agree with grammatical intuitions only for relatively short and simple sentences. The theory is an idealization, and, as before, we permit it to deviate from the intuited data in the expectation that further theory will account for the differences.

III

It might seem our job is finished. We set ourselves to giving an account of the grammarian's doings in building a grammar, and this we have done. But the reader conversant with competing accounts[7] will expect more. For, commonly, such accounts go on to talk of *linguistic theory, acquisition models, evaluation measures* and other notions related to the question of how a speaker acquires his grammar. Moreover the discussion of these notions is not a simple addition to the account of the grammarian's work in constructing a grammar. Rather it is an intrinsic part of that account. Yet why this is so is far from obvious. Constructing a theory of grammar

[7] For example, those in Noam Chomsky, "Current Issues in Linguistic Theory," in Fodor and Katz, eds., *The Structure of Language* (Englewood Cliffs, N.J.; Prentice-Hall, 1964); in Chomsky, *Aspects of the Theory of Syntax* (Cambridge, Mass.: MIT Press, 1965), ch. 1; and Katz, *The Philosophy of Language* (New York: Harper & Row, 1966).

acquisition is surely a fascinating project and one which would naturally catch a grammarian's eye. But, at first blush at least, it would seem to be a new project, largely distinct from the job of constructing grammars for individual languages. Why, then, do Chomsky and others view the study of acquisition as intrinsic to the construction of grammars for individual languages? This is the riddle that will occupy us in the present section. In the course of untangling it we will come upon some unexpected facts about grammar and its place among the sciences.

Let me begin with a puzzle. A grammar of English will generate structural descriptions for English sentences in the form of phrase markers or labeled bracketings. The labels on these brackets will be the familiar NP, VP etc. But now imagine a perverse variant of our grammar created by systematically interchanging the symbols NP and VP throughout the theory. If the change is thoroughgoing (made in all appropriate generative rules and definitions), then presumably the original theory and the variant will make exactly the same predictions about intuitions, etc. So the two would appear to be empirically indistinguishable. On what basis, then, are we to select one over the other?

To underscore the puzzle, consider a grammarian attending to the hitherto neglected tongue of some appropriately exploited and unlettered people. His grammar will likely end up generating labeled bracketings among whose labels are the familiar NP and VP. But what justification can there be for this grammar as constrasted with a variant interchanging NP and VP throughout, or yet another variant in which NP and VP are systematically replaced with a pair of symbols that occur nowhere in any grammar of English?[8]

There is a related puzzle that focuses not on the vocabulary of a grammar but on its rules. Consider any grammar or fragment of a grammar for English. With the grammar at hand it requires only modest ingenuity to produce a variant set of rules and definitions whose consequences (the entailed claims about grammaticality, grammatical relations and the rest) are identical with those of the original. Among the variants that might be produced some will differ only trivially, adding a superfluous rule perhaps, or capturing a generalization in two rules rather than one. But other variants exist which differ quite radically from the original.[9] A gram-

[8] Much the same puzzle is hinted at by Quine in "Methodological Reflections on Current Linguistic Theory," *Synthese*, XXI, 3/4 (October 1970): 386–398, pp. 390 ff.

[9] Such variants often require considerable effort to construct. Nor is it always a trivial matter to prove the equivalence of a pair of grammars.

mar is but an axiomatized theory, and it is a truism that a theory that can be axiomatized at all can be axiomatized in radically different ways. Yet each of these variants makes identical claims about the grammarian's data—not only the data on hand, but *all* the data he might acquire. They may, of course, predict incorrectly on a given point; but if one variant predicts incorrectly they all will. How then is the grammarian to decide among them?

The point of these puzzles is that grammar is afflicted with an embarrassment of riches. It is a task demanding wit and perseverance to construct a grammar that correctly captures a broad range of speakers' intuitions. Yet when the job has been done there are indefinitely many variants each of which captures the known intuitions equally well and predicts unprobed intuitions equally well (or poorly). Somehow the grammarian does come up with a single theory. What principle can he use to guide his choice?

It is in attempting to answer this question that the study of acquisition looms large in Chomsky's writings. But exactly how a theory of grammar acquisition is supposed to motivate a choice among alternative grammars is far from clear. Part of the obscurity, I suspect, stems from the fact that Chomsky, perhaps without realizing it, pursues two rather different strategies in relating the study of acquisition to the problem of choosing among alternative grammars. One of these strategies, I will contend, is thoroughly misguided and rests on a mistaken picture of what grammar is. The other is quite compatible with the account of grammar developed above and suggests an illuminating solution to the puzzles of alternative grammars. Our first project will be to dissect out these alternatives for closer inspection.

Before we begin, some terminology will be helpful. Let us call a grammar *descriptively adequate* for a given language if it correctly captures the intuitions of the speakers of the language (and the rest of the grammarian's data) within the limits of accuracy allowed by idealization. The grammarian's embarrassment of riches arises from the fact that for each descriptively adequate grammar of a language there are indefinitely many alternatives all of which are also descriptively adequate.

Now the strategy I would disparage unfolds like this:[10] When a child learns a language, he learns a descriptively adequate grammar (*dag*). He somehow "internally represents" the rules of the gram-

[10] I think this strategy is often suggested by what Chomsky says (e.g., in *Aspects of the Theory of Syntax*, pp. 24–27 and elsewhere). But my concern here is to scotch the view, not to fix the blame. So I will not bother to document details of its parentage.

mar. So if we could discover which set of rules the child has "internalized" we would be able to choose a right one from among the *dags* of the child's language. The right one is simply that grammar which the child has in fact internally represented. The study of acquisition will be designed to give us a lead on which descriptively adequate grammar the child has learned.

Let us reflect on what the child must do to acquire his grammar. The learner is exposed to what Chomsky calls *primary linguistic data (pld)* which "include examples of linguistic performance that are taken to be well formed sentences, and may include also examples designated as non-sentences, and no doubt much other information of the sort that is required for language learning, whatever this may be" (*ibid.*, p. 25). When he has succeeded in learning his language the child will have internalized a *dag*. In two rather different ways this grammar will specify more information about the language than is to be gleaned from the *pld*. First, the *pld* contain a modest sample of the grammatical sentences of the language; the grammar acquired generates all the grammatical sentences. Second, the *pld* contain little or no information about the structural descriptions of sentences and the grammatical relations among them; the grammar assigns structural descriptions to each grammatical sentence and entails all the appropriate facts about grammatical relations. Thus a theory of grammar acquisition must explain how the child can acquire and internalize a grammar that is significantly more informative about the sentences of the language than the *pld* he has been exposed to.

How might we build a theory that accounts for the child's accomplishment? What we seek is a model (or function) which, when given a complete account of the *pld* available to the child as input (or argument), will produce, as output (or value), the *dag* that the child acquires. Our problem is to design the model with sufficient structure so that it can correctly project from the limited *pld* to the full grammar of the language from which the data are drawn. What sort of information should the model contain?

Suppose it were discovered that certain features were shared by all known *dags*. If the grammars that shared the features were sufficiently numerous and diverse we might reasonably hypothesize that these features were universal among *dags* of natural language. We would, in effect, be hypothesizing that there is a restricted set of grammars that humans can in fact learn (in the normal way). Were such universal features to be found, our strategy suggests that we take account of them in our acquisition model. Since the

output of the model must be a *dag,* we would want to build our model in such a way that the possible outputs (the range of the acquisition function) each had the features that were universal to all *dags.* We would thus take the specification of universal features to define the class of *humanly possible grammars (hpgs).* The task of the acquisition model is to discover the correct grammar, the grammar of the language the child is actually exposed to, from among the humanly possible grammars.

There is great gain for the builder of an acquisition theory in discovering as rich a set of universal features as possible. For the stronger the restrictions on the *hpgs,* the smaller the class of such grammars will be. Thus the easier the task relegated to the other parts of the model. What remains for the rest of the model is to compare the *pld* with the class of *hpgs* and exclude those possible grammars which are incompatible with the data.

Now it might happen that the universal features we discover so narrow down the class of *hpgs* that only one *hpg* is compatible with the *pld.*[11] If this is commonly the case, our acquisition theory need contain only a specification of *hpgs* and a device for excluding those *hpgs* which are incompatible with the *pld.* If, however, there are several *hpgs* compatible with all the data the child has accumulated by the time acquisition is essentially complete, we will have to seek some further principle of selection. The principle, the strategy suggests, is to be found in an evaluation measure or weighting of *hpgs.* Some of the *hpgs* that are compatible with all the *pld* will still fail to be descriptively adequate for the child's language. Some of these may simply project incorrectly beyond the sample of the language available to the child. They will then classify as grammatical sequences that are not grammatical. Others, while projecting correctly, may miss the mark on structural descriptions or grammatical relations, specifying that sentences are related in ways other than the ways speakers in fact intuit them to be related. So what we seek in our evaluation measure is some ranking of *hpgs* that has the following property: when we exclude from the *hpgs* those grammars which are incompatible with the *pld,* the highest ranked of the *remaining* grammars is a descriptively adequate grammar of the language the child acquires. The acquisition model would then proceed by first eliminating those *hpgs* which are not compatible with the *pld,* then selecting from among those which remain the one that is highest ranked. The grammar selected is unique among *dags,* for it is chosen by a model that explains how a

[11] Chomsky suggests this possibility, *ibid.,* pp. 36–37.

child might go about acquiring the grammar he does acquire. It is this "explanatorily adequate" grammar which the child actually internalizes and which the linguist seeks to uncover.

A more detailed account of the strategy we are sketching might now go on to worry about how the appropriate evaluation measure could be discovered or what we can say about linguistic universals in the light of present knowledge. But this will not be our course. For I think we have said enough to see that the strategy is wholly wrongheaded. To begin, let us consider the possibility, mentioned briefly a paragraph back, that the universals so constrict the class of *hpgs* that only one *hpg* will be compatible with the *pld*. A moment's reflection will reveal that this is not a real possibility at all. For recall the pair of puzzles that initially prodded our interest in acquisition models. Each puzzle pointed to the superabundance of descriptively adequate grammars for any natural language. For every *dag* there are alternatives which are also descriptively adequate. But the linguistic universals were taken to be properties of all *dags*.[12] Thus each *dag* for every natural language will be among the *hpgs*. So if any *dag* is compatible with the *pld*, all its alternatives will be as well. And we have made no progress at selecting a single *dag* as the right one.

What is more, the hunt for an evaluation measure is of no real value in narrowing down the class of *dags*. The job that was set for the evaluation measure was not a trivial one. Given any body of *pld*, the evaluation measure had to rank as highest among the *hpgs* which are compatible with the *pld* a *dag* of the language from which the data are drawn. Finding such a measure would likely be a task of considerable difficulty. But, and this is the crucial point, once such a measure *has* been found there will be indefinitely many alternative measures which select different *dags* for the same body of *pld*. If the sub-class of *hpgs* compatible with a given body of *pld* contains *one dag* of the language of which the data are a sample, it will contain many. Thus if we can design a measure which ranks any one of these *dags* highest in the sub-class, there will be another measure which ranks a different *dag* highest.[13] But whatever justifi-

[12] It is essential that the linguistic universals be taken as the properties common to each descriptively adequate grammar of every natural language. An alternative notion that took the linguistic universals as the features common to each of the actually internalized grammars of every natural language would be useless in the present context, since our project is to discover which among the *dags* of a given language is internalized. And until we *know* which grammars are internalized we cannot discover which features are universal to such grammars.

[13] As is the case with alternative *dags*, some alternative measure functions will be trivially cooked up variants of the original. (E.g., simply select an arbitrary

cation there is for holding the *dag* selected by one measure to be the grammar actually internalized is equally justification for holding that the other is. And we are back where we started, with too many *dags* each with equal claim to be the "right one."

The second strategy for solving the problem, the strategy I would endorse, sets out in quite a different direction from the first. It does not propose to select among *dags* by finding the one actually internalized. Indeed it is compatible with (but does not entail) the view that *no* grammar is, in any illuminating sense, internally represented in the speaker's mind or brain, and that there is no good sense to be made of the notion of "internal representation." The second strategy approaches the multiplicity of *dag* as a practical problem for the working linguist. At numerous junctures a linguist may find himself with data to account for and a variety of ways of doing so. Among the alternatives, more than one will handle all the data available and will coincide in their predictions about facts as yet unrecorded. How is the linguist to choose? What the linguist seeks, according to this strategy, is not the grammar actually in the head (whatever that may mean) but some motivated way to select among *dags*.

The motivation is to be found through the study of acquisition models, though the goals of an acquisition model must be reinterpreted. If we suspend interest in which grammar is "internally represented" we need no longer demand of an acquisition model that, for a given body of *pld,* it produce as output a grammar that a learner exposed to the data would internalize. Instead, we ask only that the acquisition model have as output *some* grammar that is true of the accomplished speaker (i.e., some grammar that correctly describes the sentences acceptable to him, his intuitions about grammatical relations, etc). But let it not be thought that this is a trivial task. Such a model would be able to specify a grammar true of the speaker given only the (relatively scant) primary linguistic data to which the speaker was exposed. To do this would be a monumentally impressive feat realizable, for the foreseeable future, only in linguistic science fiction.

How can such a model be built? In attending to the more demanding model of the first strategy, our first move was to linguistic universals, the properties shared by all *dags*. The analogous role in the present strategy can be played by properties less difficult

---

dag of the language from which the *pld* is drawn and place it highest under the evaluation measure, leaving the rest of the measure unchanged.) Others will exist which differ from the original in more substantial ways.

to discover. For suppose we have a single descriptively adequate grammar of a particular natural language. Might it not be reasonable to take as many properties of that grammar as possible as "quasi-universals"? "Quasi-universal" properties play just the role that universals did in the first strategy—they constrain the output of the acquisition model. The quasi-universals, then, define a class of "quasi-humanly possible grammars" which are the only possible outputs of the acquisition model. The terminology is adopted to stress the parallel with the first strategy. But there are important differences. For quasi-universals are in no sense universals—there is no claim that all *dags* must share them. Nor does the class of quasi-humanly possible grammars pretend to exhaust the class of grammars that humans can learn;[14] it simply coincides with the possible outputs of the acquisition model.

As was the case at the analogous point in the first strategy, there is profit in taking the quasi-universals to be as strong as we can. For the stronger the quasi-universals, the smaller the class of quasi-*hpgs* and thus the easier the task that remains for the rest of the model. Indeed, it would not be unreasonable as a first guess to take *all* the properties of the single *dag* as quasi-universals.[15] But this clearly will not do. For then the output class of the acquisition model would have but a single member. Rather, our principle in deciding whether to take features of our single *dag* as quasi-universal is this: take as quasi-universal as many features of the *dag* as possible, provided only that the resultant class of quasi-*hpgs* contains at least one quasi-*hpg* for each natural language. The remainder of the model will contain (at least) a component testing the compatibility of quasi-*hpgs* with the accumulated *pld*. Note that, on this second strategy, it is indeed possible that the quasi-universals so narrow down the class of quasi-*hpgs* that only one *hpg* will be compatible with any given body of *pld*. If this is the case, then a specification of the quasi-universals and a compatibility-testing device of the sort lately considered would complete an acquisition model. But if we cannot discover quasi-uni-

[14] Indeed, if we abandon the notion of internal representation, it is no longer clear that it makes sense to speak of a child "learning" a grammar. When the child succeeds in mastering his mother tongue, each *dag* of that tongue is true of him. But he surely has not learned *all* these *dags*. What, then, is the "cash value" of the claim that he has learned any one of them?

[15] During the John Locke Lectures at Oxford in 1969, Chomsky suggested that were a Martian linguist to come to earth in the midst of an English-speaking community, his most reasonable first hypothesis would be that the ability to speak English is entirely innate. I suspect that Chomsky's remark and the present observation are directed at basically the same point.

versals of this strength, we will again resort to an evaluation measure. As with the first strategy, what we seek is a ranking of quasi-*hpgs* which, when we exclude from the quasi-*hpgs* those grammars incompatible with a given body of *pld,* ranks highest among the remaining quasi-*hpgs* a grammar that is descriptively adequate for the language from which the *pld* was drawn. Since we are making no claim that the selected grammar is "actually internalized" we need not be concerned that there may be several such evaluation measures. Our project is the highly nontrivial project of producing a model that takes *pld* as input and yields an appropriate *dag* as output. *Any* evaluation measure that does the trick will be suitable.

The outline we have given of the construction of an acquisition model is, in a crucial respect, misleading. For it suggests that the model builder is bound irrevocably by the first *dag* he constructs. He takes as quasi-universal as many properties of this grammar as he can get away with, weakening the quasi-universals only when he comes upon some language no *dag* of which would be included among the quasi-*hpgs* if the stronger quasi-universals are retained. Actually, of course, matters are much more flexible. There is room for substantial feedback in both directions as work proceeds on the model and on individual grammars. The overriding concern is to make both the individual grammars and the acquisition model as simple and as powerful as possible. If at a given juncture it is found that adhering to the working hypothesis about the acquisition model will substantially complicate construction of grammars for one or more languages, he will try to alter the model, even if this may require altering or abandoning the original grammar from which the earliest hypothesis about quasi-universals was drawn. And, on the other side, if in constructing a particular *dag* a certain choice of how to proceed would accord well with the working hypothesis about the acquisition model, then he will be inclined to make that choice even if the resulting grammar is somewhat less elegant than another which would result from an alternative choice. There is no circularity here, or at least, to crib a phrase, the circularity is virtuous. Through this process of mutual adjustment progress on the acquisition model and on particular grammars can take place simultaneously.

Notice, now, that the strategy we have been detailing will solve the puzzles with which we began. An acquisition model provides motivation for selecting one *dag* over another, though both do equally well at predicting intuitions and such. The grammar to be chosen is that which accords with the quasi-universals. And, if

several do, the grammar chosen is the one the evaluation measure ranks highest. Thus the grammar chosen will be preferred to its descriptively adequate competitors because it is more closely parallel to successful grammars for other languages and integrates more successfully into a model of grammar acquisition.

The account we have given of the second strategy has the further virtue of according well with actual linguistic practice. It is simply not the case that, when speculating about "linguistic universals," Chomsky and his followers set out to survey a broad range of languages and collect those features common to all the grammars. Rather, speculation is based on the study of a single language, or at best a few closely related languages. A feature of a grammar will be tentatively taken as "universal" if it is sufficiently abstract (or nonidiosyncratic) to make it plausible that the feature could be readily incorporated into a grammar of every natural language. If "universals" are taken to be features common to all *dags,* this speculation about universals would be quite mad. But in the light of the second strategy the speculation appears as a thoroughly reasonable way to proceed.

An element of indeterminacy still lurks in our second strategy. And if I am right in identifying this strategy with the generative grammarian's practice, then the indeterminacy infuses his theory as well. In constructing an acquisition model, the first few plausible (approximations of) descriptively adequate grammars have a profound influence. For it is the abstract features of these grammars which are taken as quasi-universals. Yet the selection of these first *dags* over indefinitely many alternatives is completely unmotivated by any linguistic evidence. Which *dag* is first constructed is largely a matter of historical accident. But the accident casts its shadow over all future work. The acquisition model serves to direct future research into the channel forged by these first grammars, even though there are indefinitely many other possible channels available. Nor does the flexibility we stressed three paragraphs back eliminate the indeterminacy. There we noted that, if an original choice of quasi-universals led to overwhelming difficulties in constructing a grammar for some previously neglected language, the universals might be patched and the early grammars that suggested them might be abandoned. But the new choice of quasi-universals has no more claim to uniqueness than the old. For they too will be abstracted from *dags* that were selected over competitors largely by virtue of historical accident.

To the appropriately conditioned reader this indeterminacy

will appear familiar enough. It bears strong analogy with Quine's thesis of the indeterminacy of translation.[16] Quine's analytical hypotheses, like the first *dags,* are underdetermined by the data. The selection of one *dag* or one set of analytical hypotheses is largely a matter of cultural bias or historical accident. But once a *dag* or a set of analytical hypotheses has been formulated, it has profound effects on the remainder of the translation theory (for analytical hypotheses), or on the acquisition model and *dags* for other languages. Both analytical hypotheses and early *dags* are susceptible to later tampering; but neither a patched *dag* nor a patched analytical hypothesis has any more claim to uniqueness than the originals.

My departure from Quine comes on the score of the *implications* of the indeterminacy. Were Quine to grant that grammars and translation manuals share a sort of indeterminacy,[17] he would presumably conclude that for grammars, as for translations, modulo the indeterminacy, there is nothing to be right about. On this view there is no saying that one *dag* of a language is more correct than another, except relative to a given set of quasi-universals. Yet the selection of quasi-universals, like the selection of analytical hypotheses, is in part quite arbitrary. My dissent comes in the step that passes from recognition of arbitrariness in quasi-universals or analytical hypotheses to the claim that there is (modulo the indeterminacy) nothing to be right about. For I think that, *pace* Quine, the same indeterminacy could be shown lurking in the foundations of every empirical science. Grammar and translation are not to be distinguished, in this quarter, from psychology or biology or physics. If we are disinclined to say that in all science, modulo the indeterminacy, there is nothing to be right about, it is because the theories we are willing to allow as correct are those whose arbitrary features have the sanction of tradition. But all this is to stake out my dissent, not to defend it. The defense is a project I must postpone until another occasion.

IV

Our sketch of the grammarian's doings is all but complete. We have surveyed the data to which he attends and indicated the nature of

[16] Cf. "Speaking of Objects," *Proceedings and Addresses of the American Philosophical Association,* XXXI (1957/8): 5–22; "Meaning and Translation," in Fodor and Katz, *The Structure of Language, op. cit.; Word and Object* (Cambridge, Mass.: MIT Press, 1960), ch. II; and *"Ontological Relativity,"* this JOURNAL, LXV, 7 (April 4, 1968): 185–212, reprinted in *Ontological Relativity, and Other Essays* (New York: Columbia, 1969).

[17] There is evidence that he would. Cf. "Methodological Reflections . . . ," *op. cit.*

the theory he builds upon his data. It remains to say something of the interest of the grammarian's theory and to set out the relation between his theory and the speakers whose intuitions and behavior are his data.

As I have depicted it, a grammar is a modest portion of a psychological theory about the speaker. It describes certain language-specific facts: facts about the acceptability of expressions to speakers and facts about an ability or capacity speakers have for judging and classifying expressions as having or lacking grammatical properties and relations.

The modesty of a grammar, on my account, stands in stark contrast to more flamboyant portraits. On Jerrold Katz's view, a grammar is a theory in physiological psychology whose components are strongly isomorphic to the fine structure of the brain. "The linguistic description and the procedures of sentence production and recognition," according to Katz, "must correspond to independent mechanisms in the brain. Componential distinctions between the syntactic, phonological, and semantic components must rest on relevant differences between three neural submechanisms of the mechanism which stores the linguistic description. The rules of each component must have their psychological reality in the input-output operations of the computing machinery of this mechanism."[18] Though Katz's claims about grammar are more expansive than those I have made, the evidence he uses to confirm a grammar is of a piece with the evidence indicated in my account. Thus it remains something of a mystery how the grammarian has learned as much as Katz would have him know about the structure of the brain, having left the skulls of his subjects intact.

Less imaginative than Katz's view, but still not so sparse as mine, is a story about grammar put forward by Chomsky.[19] On this account a grammar describes the speaker's "competence"—his knowledge of his language. The speaker is held to have a large and complex fund of knowledge of the rules of his grammar. The grammarian's theory mirrors or describes the knowledge that the speaker has "internalized" and "internally represented." Chomsky's view is intriguing, though an explicit unpacking of the metaphors of "internalization," "representation," and the rest can prove an exasperating task. My own view is that the notion of competence is explanatorily vacuous and that attributing knowledge of a grammar to a speaker is little more plausible than attributing knowledge of

[18] "Mentalism in Linguistics," *Language*, XL, 2 (April/June 1964): 124–137, p. 133.

[19] In *Aspects of the Theory of Syntax, op. cit.*, and elsewhere.

the laws of physics to a projectile whose behavior they predict. But the issues are complex, and I have aired my views at length elsewhere.[20] I will not rehash them here. What is important to our present project is the observation that, on the account of grammar and acquisition models we have constructed, no knowledge claim is *needed.* A grammar is a theory describing the facts of acceptability and intuition; a grammar-acquisition model is a theory specifying a grammar which comes to be true of a child, as a function of the linguistic environment in which he is placed. Grammar and the theory of grammar acquisition are bits of psychological theory.

If our account of the grammarian's activity is accurate, then it is perhaps misleading to describe him as constructing a theory of the language of his subjects. Rather he is building a description of the facts of acceptability and linguistic intuition. A theory of a language seriously worthy of the name would provide some insight into what it is to *understand* a sentence, how sentences can be used to communicate and to deal more effectively with the world, and into a host of related questions that we have yet to learn to ask in illuminating ways. But a grammar does none of this. Indeed, it is logically possible that there be a person whose linguistic intuitions matched up near enough with our own, but who could neither speak nor understand English. Such a person would serve almost as well as an English speaker as an informant for constructing a grammar of English, provided only that we shared a metalanguage in which we could question him about the sequences of sounds he did not understand. What is important about this bit of fiction is that it is *only* fiction. It is an empirical fact that comprehension and intuition run in tandem. And this fact provides the beginning of the answer to a question that will likely have begun to trouble the reader: Of what interest is a grammar? If a grammar is not, in any exciting sense, a theory of a language, why bother constructing it?

The answer is twofold. First, there is substantial correspondence between the grammatical sentences and the sentences we do in fact use for thought and communication; grammatically related sentences are understood in similar ways[21] (though in our present state of ignorance we have no serious understanding of what it is to

[20] "What Every Speaker Knows," *Philosophical Review,* LXXX, 4 (October 1971): 476–496, and "What Every Grammar Does," to appear in *Philosophia.*

[21] Cf. Chomsky *Syntactic Structures* (The Hague: Mouton, 1957), p. 86: "the sentences (i) *John played tennis* [and] (ii) *my friend likes music* are quite distinct on phonemic and morphemic levels. But on the level of phrase structure they are both represented as *NP-Verb-NP; correspondingly, it is evident that in some sense they are similarly understood.* (Last emphasis added.)

"understand sentences in similar ways"); the ability to speak and understand a language is an empirically necessary condition for the possession of linguistic intuitions about the expressions of the language. So one reason for studying grammar is the hope that these overlaps and correlations can be exploited to yield deeper insight into the exciting phenomena of comprehension and communication. Once we have the sort of description of acceptability and linguistic intuition provided by a grammar we can begin to seek an explanation of these facts. We can ask what psychological mechanisms underlie the speaker's ability to judge and relate sentences as he does. The parallels between linguistic intuition and other language-related phenomena make it reasonable to hope that insight into the mechanisms underlying intuition will explain much else about language as well. But hope is not to be confused with accomplishment. If we fail to recognize how modest a theory a grammar is, we can expect only to obscure the extent of our ignorance about language, communication, and understanding.

A second reason for doing grammar is that it is something to do. In grammar, at least, we have a coherent set of data that we know how to study, intelligible questions to ask, and some clear indication as to how we can go about answering them. Acceptability and grammatical intuitions are language-related phenomena about which we have the beginnings of an empirical theory. Few other approaches to the phenomena of natural language fare as well. Thus grammar is a natural focus of attention for the investigator concerned with language. It is an entering wedge to a theory of a language, and, for the present at least, there are few competitors.

STEPHEN P. STICH

The University of Michigan

INDETERMINACY OF TRANSLATION

THIS PAPER is composed of four parts. In part I, I present and clarify Quine's thesis of indeterminacy of translation. In part II, I argue that Quine's claim of indeterminacy of translation is not adequately supported because it rests on the unwarranted assumption that acceptance and rejection of theoretical sentences is not empirically determined to the degree that acceptance and rejection of nontheoretical sentences is determined. In part III, I reject Quine paradigm of translation because it wrongly supposes that a language is a set of dispositions to behave in certain