

4

PSYCHOLOGICAL PROCESSES AS LINGUISTIC EXPLANATION¹

Herbert H. Clark and Susan E. Haviland

*Department of Psychology
Stanford University*

In U. Cohen (Ed.), *Explaining
Linguistic Phenomena.*
Washington, D.C.: Hemisphere
Publishing Corp., 1974

In the halcyon days of transformational theory, many linguists, following Chomsky's lead, were confident that a sizeable chunk of English could eventually be accounted for with self-contained grammars capable of generating all, and only, the "grammatical" sentences of the language. A strong premise behind this early work was that there exist linguistic structures that can be studied independently of such psychological factors as limited memory capacity, world knowledge, belief systems, and processing capabilities. Accordingly, much use was made of a construct which dates back at least to Frege, the "ideal speaker-listener," whose competence or linguistic knowledge can be described independently of his psychologically based performance factors. Little by little, however, it has become apparent that the borderline between the purely linguistic and the psychological aspects of language is fuzzier than was once supposed. In fact, it may not exist at all.

There are three main reasons why many investigators of language have come to question the competence-performance distinction. First, several

¹The preparation of this paper was supported in part by a grant from the National Institute of Mental Health MH-20021. We thank Eve V. Clark for her comments on previous drafts of this paper. We also express our appreciation to W. C. Watt, whose perspicacious and persuasive discussions on many of the points touched on in this paper have had a profound influence on the general position expressed here. He should not, of course, be held responsible for anything we say here; his own views appear elsewhere (see especially Watt, 1970, 1972, 1973a, 1973b).

traditionally psychological factors—e.g., processing limitations, a priori beliefs, and communicative intentions—have by now become quite important in linguistic descriptions of a variety of phenomena. At least some linguists (e.g., Gordon and Lakoff, 1971; R. Lakoff, 1972) have argued that these factors cannot legitimately be divorced from the grammar of English. Second, many linguists have come to find the study of linguistic competence, as defined by the early transformational grammarians, too confining. There are simply too many interesting linguistic phenomena that do not fit under this rather small umbrella, yet are amenable to linguistic investigation. Thus, linguists have begun to apply their skills to such problems as the structure of conversations (Gordon & Lakoff, 1971; Heringer, 1971; Morgan, 1972), the relation of meaning to context (Geis & Zwicky, 1971; R. Lakoff, 1972), the perceptual difficulty of surface structures (Bever, 1970; Bever & Langendoen, 1971; Kimball, 1973), the production of speech errors and speech hesitations (Fromkin, 1971, 1972; James, 1972), and other similar phenomena. In the examination of such phenomena it has become apparent that many performance factors appear to be governed by rules and based on principles that may be universal. This has made it all the more difficult in principle to distinguish performance factors from competence factors. Third, linguists—beginning most forcefully with Chomsky (e.g., 1965)—have set out to psychologize about language, asking such questions as: Is linguistic competence innate? How does it manifest itself in language acquisition? Is man unique in this linguistic competence? Can one explain historical change by referring to the process of language acquisition? etc. This sort of psychologizing—some of it very traditional, in fact—has led to various conjectures about the psychological basis for the form and content of linguistic rules, and this has necessarily involved reference to performance factors. Bever and Langendoen (1971), for example, have argued that certain transformations exist principally to facilitate the perceptual processing of sentences, and Kimball (1973) has made similar conjectures. In short, many linguists have felt obliged to deal with at least the interface between performance and competence.

Traditionally, linguists have attempted to account for language with models that are essentially static. Such models have been designed to characterize what people have to know in order to produce and comprehend utterances, not what people actually do in such production and comprehension. Given that linguists have generally limited themselves to such models it is hardly surprising that they have wanted to maintain the competence-performance distinction. Whereas competence seems amenable to explication by such static models, performance does not. But if the linguist wants to

introduce into his study of language such factors as processing limitations, belief systems, and communicative intent, he may well have to abandon such static models because of their inherent inability to deal with mental states that change from moment to moment. For this and other reasons, linguists may be forced to turn to the type of process models cognitive psychologists typically work with.

Process models differ from the typical linguistic model mainly in that they are concerned with mental processes that occur in time. In the static linguistic model, of course, there are rules that develop the left-to-right ordering of the formatives in surface structure, but beyond this, time is not involved. Indeed, time seems to enter this model by the back door. Note that in linguistic theories the sequencing is of the superficial elements in the output of speech, not of any internal process required in the construction of this sequence. So although it does not order mental events at any previous stage and it has no other time based processes, the grammar, as if by magic, produces as its final step a sequence of elements ordered in time. In process models, on the other hand, time and the sequencing of mental events are central to the whole enterprise. The typical process model is stated in terms of a series of mental states, each represented in some symbolic form, and a collection of rules that change one mental state into the next. These mental states are assumed to occur in a particular sequence—although sequences can occur in parallel—and each change of state requires a certain increment of time to perform. (For examples of simple process models, see Sternberg, 1969; Clark & Chase, 1972; Garrod & Trabasso, 1973). Because of their inherent properties, these models are potentially able to accommodate: (1) the moment to moment changes in the mental states of the speaker or listener; (2) the listener or speaker's knowledge of language, as represented in the rules required for changing from one mental state to the next; (3) the speaker or listener's processing capacity, as found in the limitations on what rules can be applied, how many can be applied at one time, and how long each rule takes; and (4) the time required to say or understand an utterance. In short, process models, unlike static models, are not forced to make the strong assumption that one's knowledge of language can be separated from the role this knowledge plays in the actual comprehension or production of utterances. They are able to accommodate both competence and performance, as traditionally defined, simultaneously.

In what follows we would like to explore the possibility of introducing process models as accounts for certain linguistic phenomena. We will first look at some specific linguistic phenomena that seem to require something more than a static description for their explanation. After that we will

examine one particular processing strategy in more detail to see how useful it is in accounting for some facts about thematic structure in English. Finally, we will collect our separate points in a rather general proposal for the study of natural language as a part of the study of the mind.

FOUR LINGUISTIC EXAMPLES

We will first take up several examples that suggest that linguists may have to turn to process models to account for linguistic phenomena that are well within the province of the traditional linguist.

Example 1

As our first example, we consider the problem of anaphoric reference into more or less penetrable anaphoric islands. As Watt (1973a, 1973b) has shown recently, pronouns vary widely in their ability to refer to their antecedents (anaphors) contained within other constructions. Consider the sentences in (1):

- a. Because it's such a great city, I've always wanted to be a San Franciscan.
 - b. I've always wanted to be a San Franciscan, because it's such a great city.
- (1)

In (1) the pronoun *it* is meant to refer to *San Francisco*, which in turn is contained within the surface word *San Franciscan*. The question arises whether *it* can penetrate *San Franciscan* to get at its anaphor *San Francisco*. Watt has shown that there are a number of factors affecting whether such pronouns are successful or not. First, the pronoun is typically more successful in its penetration when it precedes its anaphor, as in (1a), than when it does not, as in (1b). That is, (1a) seems more acceptable than (1b) despite the fact that they are identical except for the preposing of the subordinate clause in (1a), an operation that normally has no effect on pronominalization in the subordinate clause. Second, pronouns are typically better able to penetrate to their anaphors the more specific the pronoun is, semantically, in referring to the anaphor. Consider (2):

- a. I met a San Franciscan who just loved it there.
 - b. I met a San Franciscan who just loved it.
- (2)

The *there* in (2a) is more specifically locative than the *it* in (2b), and so (2a)

sounds better—much better—than (2b). Third, contrastive stress of the anaphor can also lead to easier penetration, as illustrated in (3):

- a. Before Vera could NAIL the sign to the mast, Cornell did it with a staple.
- b. Before Vera could nail the sign to the mast, Cornell did it with a staple. (3)

In these examples, the *did it* can penetrate *nail* (= “fasten with nails”) to get to “fasten” more easily in (3a) with the stressed NAIL than in (3b) with the unstressed *nail*. The penetration of words can even occur when the pronoun has been deleted, as in (4):

The Dutch are very tolerant toward other nations. (4)

Here, *other nations* mean nations other than it (= Holland), where *it* has been deleted. Nevertheless, the missing *it* is able to penetrate *Dutch* (= “residents of Holland”) to get to Holland. Indeed, it appears that such penetration is often more acceptable when the pronoun has been deleted than when it has not. Consider (5):

- a. I met a Dutchman who was so set in his ways that he wouldn't live anywhere else.
- b. I met a Dutchman who was so set in his ways that he wouldn't live anywhere but there. (5)

Although (5a) and (5b) are virtual paraphrases, (5a) seems more acceptable despite the fact that the *there* in *anywhere but there* has been deleted. Note that in (4) and (5), the *it* and *there* are penetrating *Dutch* and *Dutchman* to get to *Holland*, or *The Netherlands*, which are phonologically unrelated to the words *Dutch* and *Dutchman*.

The traditional linguist might want to relegate these phenomena immediately to a theory of performance, but this seems unwise for two reasons. First, these phenomena bear on whether the grammar should contain, say, *Dutchman* as an undecomposable lexical item, or as a lexical item that decomposes into something like “resident of Holland,” the elements of which are now available for anaphora by the rules of the grammar. Note that there is no simple answer to this question. Apparently, given the evidence Watt has laid out, the elements of the decomposed *Dutchman* must be available for

anaphora, but the degree to which they *are* available in any particular case varies with all sorts of conditions. The word *Dutchman* behaves like the phrase *resident of Holland* in some instances, but not in others: whereas *resident of Holland* is always penetrable, *Dutchman* is not. Second, the penetrability of some anaphoric islands appears to be affected simply by the passage of time, a factor no competence theory has been willing to include. Consider (6):

- a. I've been a life-long San Franciscan. It's such a beautiful city.
 b. I've been a life-long San Franciscan. (Why?) It's such a beautiful city. (6)

As Watt has pointed out, the passage of time alone often makes a pronoun sound better. Note that (6b) seems slightly better than (6a) because of the extra time between *San Franciscan* and *it*. It is as if this extra time allows the listener time to extract *San Francisco* from *San Franciscan* so that he will have the appropriate anaphor by the time the pronoun *it* appears. (In the extreme, of course, passage of time will eventually make a pronoun seem worse, since the listener will have forgotten the previous sentence and the possible anaphors it contains.)

The facts laid out by Watt suggest quite a different model for the handling of these complex nouns, a model, moreover, that posits mental states and changes in mental states. Consider the word *Dutchman*. Part of what one knows about this word is that it can be decomposed into the elements "resident of Holland," but this decomposition does not invariably take place, and when it does take place, it takes time. Now it is rather straightforward to encompass most of the facts laid out by Watt. First, when the pronoun precedes the anaphor, the listener is "looking for" an anaphor, and so he is more likely to decompose the word *Dutchman* into its elements in search of the anaphor than otherwise. Second, when the pronoun is more specific semantically, the listener is more likely to decompose *Dutchman* in search of the anaphor, since he can rule out other possible anaphors more readily. Third, contrastive stress on the anaphor, as in (3a) above, signals the advent of a contrasting element, and so when the pronoun comes later, the listener is ready to consider the anaphor and its decomposition. Watt has further shown that only some elements in the decomposed anaphor can be contrasted, and this suggests that contrastive stress on the anaphor may force its decomposition in order to expose the contrasted element. Finally, the anaphor *Dutchman* is more likely to be decomposed the more time is allowed, since the decomposition itself requires time. This would account for the phenomenon

illustrated in (6). Note that such a model—suggested in almost this form by Watt—differs quite drastically from the typical competence model, which would assume that if *Dutchman* is decomposable, it is available for anaphora regardless of these sequential and temporal factors. Although this particular process model may not be correct, it does show that process models are particularly well suited to account for this type of linguistic phenomenon.

Example 2

For our second example, we take up the problem of unique pronominal reference. Consider sentences (7a) and (7b), as recently discussed by Lakoff (1973), among others:

- a. John and Mary entered the room, and he took off his coat. (7)
- b. *John and Bill entered the room, and he took off his coat.

The *he* in (7a) is unambiguous in that John is the only male in the previous clause that *he* could refer to. In (7b) *he* could refer to either John or Bill, and so (7b) is a bad sentence. Lakoff has proposed that the unacceptability of (7b) be accounted for by a transderivational constraint. The idea, roughly, is this. Whenever the surface structure resulting from a given derivation is identical to that of another in which the surface pronoun refers to a different noun, then, both derivations are blocked.

Regardless of the merits of such a transderivational constraint, there is one rather strong objection to it as it stands now. The problem is that it is defined only over linguistic entities—identical surface structures from different derivations—whereas the issue it was meant to solve is much broader than that. Consider (8) and (9):

- a. See John and Mary over there? He has rabies.
- b. See those two people over there? (One is male, the other female.) He has rabies. (8)
- a. *See John and Bill over there? He has rabies.
- b. *See those two people over there? (Both are male.) He has rabies. (9)

While (8a) and (8b) are acceptable, (9a) and (9b) are not. As presently formulated, however, this transderivational constraint accounts for (8a) and (9a), but not for (8b) and (9b), which reflect the identical phenomenon. In

the latter two instances, whether the pronoun *he* is ambiguous or not depends entirely on the listener's knowledge of the two people being referred to by *those two people over there* in the first sentence. The sentences in question need not even contain referring expressions, making the transderivational constraint as now formulated completely inapplicable. Note that *Look over there—he has rabies* will be considered acceptable to a listener only if he has been directed by the first sentence to a unique male to which *he* could refer. This example alludes to the much more general problem of how the grammar is to account for the relation between deictic and nondeictic pronominal reference.

The examples in (8) and (9), therefore, suggest that a more general theory of transderivational constraints should refer not to linguistic entities at all, but rather to the states of knowledge of the speaker and listener at the moment of utterance. In (8) the listener can assign a unique reference to *he* either from linguistic context or from visual inspection; in (9), however, he cannot. To capture this generalization, an appropriate transderivation constraint will have to be able to refer to the knowledge derived from extralinguistic sources in the same terms as it does that derived from linguistic context. This requires a process model for several reasons. First, non-linguistic knowledge is not necessarily a fixed structure that can be referred to in the same way as linguistic structures. Non-linguistic knowledge varies from moment to moment, requiring an account of how the listener constructs such states of knowledge, how he consults this knowledge in the application of the rule, and how he eliminated ambiguity on the basis of such consultations. Second, and more significantly perhaps, the linguist must face the issue of how non-linguistic knowledge is to be represented. As typically written at the present time, linguistic rules range only over P-markers and other similar "linguistic" structures. If the linguist is to retain this assumption, he will have to maintain that non-linguistic knowledge is first converted into such "linguistic" structures for reference by the linguistic rules. This is a rather strong assumption, and it seems rather unlikely at the outset. An alternative route is to abandon the notion that linguistic rules must refer only to things "linguistic" and to accept the proposal, put forward by a number of linguists and psychologists by now, that there are cognitive structures, available for such rules, that are neutral with respect to the various subsystems making up the complete mental apparatus. It matters little which route the linguist takes, however, for he will be unable to avoid making significant claims about psychological matters that have traditionally lain outside the province of linguistics. These matters will include claims about how people construct,

change, consult, and abandon states of knowledge gained from non-linguistic sources, claims that clearly require some sort of process model.

Example 3

Several linguistic proposals already make reference to the listener's knowledge—in particular, Geis and Zwicky's (1971) important discussion of "invited inferences" and Gordon and Lakoff's (1971) provocative proposal about "conversational postulates." As an example of invited inference, consider the following sentence:

John'll die soon if he has rabies. (10)

Most people construe this sentence to mean "John will die soon *if and only if* he has rabies," even though (10) does not necessarily entail that John won't die soon if he doesn't have rabies, the *only-if* part of this interpretation. People tend to "perfect the conditional" in instances like this. The interesting problem here, as Geis and Zwicky point out, is how to account for which conditionals people will perfect and which they will not. It appears that such an account will have to refer to the knowledge or beliefs of the listener at the moment of utterance. Typically, conditional perfection seems to arise because of a rule of conversation Grice (1967) has called Quantity, which itself is a subdivision of the more general Cooperation Principle. The rule is this: "Make your contribution as informative as is required (for the current purposes of the exchange)." Sentence (10) would hardly be informative if the speaker already knew that John was going to die soon no matter whether or not he had rabies. If this were the case, the speaker should have said *John is going to die soon*, an assertion not conditionalized on whether John has rabies. It follows, then, that the listener, who assumes that the speaker is obeying the rule of Quantity, can immediately infer that the speaker believes that John will not die soon if he doesn't have rabies. One general purpose of the rules of conversation is to avert being misunderstood by the listener. To do this the speaker must tell all he knows that is relevant about the topic at hand, not just part of it. So for the speaker to say

Some of John's children are girls. (11)

when he knows that all of them are girls would be not to tell all he knows, even though the sentence is, technically speaking, truthful. In short, to account for invited inferences, as Geis and Zwicky imply, linguists will have to

be able to refer to the momentary beliefs or knowledge of the listener, including such rules of conversation as those proposed by Grice.

The conversational phenomena that have been discussed recently by Gordon and Lakoff (1971), Heringer (1971), Morgan (1972), and Sadock (1970, 1972) have much the same character as invited inferences. Consider (12):

Can you open the door? (12)

As these investigators have noted, (12) has the superficial form of a question and will sometimes be interpreted as a question about the possibility of the speaker opening the door. In many contexts, however, this sentence will be taken as a request to open the door. How could this be? There are various explanations proposed in this literature, and we will assume Gordon and Lakoff's for convenience here. Their data is that the listener will interpret (12) as a request based on a combination of three things: (1) the literal interpretation of (12) as a question about the possibility of the addressee's opening the door; (2) the listener's knowledge of the situation that implies that the speaker knows very well that that act is possible; and (3) a so-called conversational postulate which says, roughly, that when the speaker queries the possibility of an act by the addressee which both know to be obvious, the speaker is conveying a request. Such conversationally conveyed requests can vary from the very obvious to the very indirect:

- a. Would you please open the door?
 - b. Is it possible for you to open the door?
 - c. I wouldn't mind if you opened the door.
 - d. It's hot in here.
- (13)

In any case, to account for many of these requests it appears necessary to make reference to the listener's knowledge of the situation as well as to his knowledge of the structure of conversations, as contained in the conversational postulates.

Since such knowledge changes from moment to moment, the ultimate description of conversationally conveyed meaning will require a dynamic theory that makes reference to the construction of knowledge states, the consultation of such states, and the changing of these states, all of which, of course, occur in real time. Consider, for example, the sentences in (14):

- a. Will you open the door? And if so, why?
b. Will you open the door? Please? (14)

Note that *Will you open the door?* can be construed as one of two things, a question about what the speaker predicts will happen, or a request to open the door. After hearing this much of (14a) or (14b), then, the listener could have either one of these two readings. But on encountering the second sentence in both sequences, the first sentences become disambiguated. Now (14a) must be a question, and (14b) a request. An important issue here is, what is going on in the process of disambiguation? There are at least two possible models. First, the two readings can be considered on an equal footing. *Will you open the door?* is ambiguous just like *Flying planes can be dangerous* in that neither reading is logically prior to the other. In this case, the second sentences in (14a) and (14b) simply force the listener to abandon the wrong reading, if he had originally chosen it, and to choose the correct one. Second, the question reading could be logically prior to, or included in, the request reading, as proposed by Gordon and Lakoff. In this instance, the second sentences in (14a) and (14b) lead to different processes. If the listener has taken the question reading for (14b), the *please* will force him to bring conversational postulates to bear and will lead him to construct the request reading. On the other hand, if the listener has taken the request reading for (14a), the *why* will force him to abandon the conversational postulates previously brought to bear in the interpretation, and this will leave him the logically prior question reading. Thus, these two models differ in whether the disambiguation is one of changing from one reading to another previously unconsidered reading, or one of adding or subtracting additional information onto a basic reading. Although we do not know how to test these models at this time, our point is this: The linguistic question of how the two readings of *Will you open the door?* are related is fundamentally inseparable from the psychological question of how a process model is to represent the process of disambiguation. If disambiguation can be shown to occur by the second method given above, then it would be inconsistent for a linguistic theory to assume two readings as found in the first model. Indeed, it would seem that the psychological question is the more fundamental one whose answer should properly dictate the answer to the other.

Example 4

As our final example, we would like to consider the problem of productivity in the interpretation of certain lexical items. Consider the

so-called denominal verbs, which have recently been discussed by Green (1972), McCawley (1971), and Watt (1973a, 1973b):

- a. John hammered the nail into the wall.
- b. John nailed picture to the wall. (15)
- c. John trucked the goods from Denver to San Francisco.

The verbs *hammer*, *nail*, and *truck* in (15) appear to be systematically related to their respective noun counterparts, such that they have the interpretations shown in (16):

- a. John drove the nail into the wall with a hammer.
- b. John fastened the picture to the wall with a nail. (16)
- c. John transported the goods from Denver to San Francisco by truck.

An important question here is whether or not one could derive the verbs from the nouns in the present grammar of English. The relation between the *hammers* of (15a) and (16a), for example, may be of historical interest, but have no bearing on the way they are understood in present day English.

The question, of course, is a complex one, and we do not pretend to have the answer to it. But several considerations lead us to suggest that at least some denominal verbs must be derived by rule, and furthermore, this rule must make reference to some rather intricate knowledge the listener has of the objects to which the nouns contained in the verbs refer. Consider, for example, the sentences in (17):

- a. Wimsey cloakroomed his hat.
- b. Nixon cloakroomed each of the senators into voting for his bill. (17)

It is unlikely that most American speakers will have heard either of these two uses of *cloakroom* before, yet most will be able to understand them quite readily. The first, adapted from Dorothy Sayers, will be interpreted as in (18a), and the second, adapted from a sentence heard in Washington, as in (18b):

- a. Wimsey put his hat in the cloakroom.
- b. Nixon used high pressure techniques in the cloakroom to persuade each of the senators to vote for his bill. (18)

Where do these interpretations come from if people have never heard the verb *cloakroom* before? And how could they differ in these two instances? A rough answer seems to be that the listener relies on his knowledge of the normal use of the objects referred to by the noun *cloakroom*. Typically, cloakrooms are used for the temporary storage of coats and hats. (17a) gets its interpretation by changing this normal use into a verb—something like “store in a cloakroom”—and then inserting this interpretation for the verb *cloakroomed*. Actually, (17a) is ambiguous between a punctual and durative reading of *cloakroomed*—Wimsey either did it at six o'clock or did it for one hour—and this too would have to follow somehow from a productive rule for the interpretation of *cloakroom*. In Washington, unlike London, cloakrooms have another normal use, and that is as a place for giving private high pressure sales talks to members of Congress. This use turns into a rather complicated, but nevertheless well understood verb, something like “use the cloakroom as a place to persuade with high pressure salesmanship.” When substituted for the verb *cloakroomed* in (17), this paraphrase produces the reading most informed Washingtonians would attribute to the sentence.

The verb *cloakroomed* is not an isolated example. One can choose any number of nouns in English that have an agreed upon normal use and verbize them with greater or lesser success. As yet it is difficult to see just why some nouns work and others do not, but our guess is that they will be found to fall into a regular pattern. The most important point, however, is that the linguistic account for such denominal verbs apparently has to make reference to knowledge of the objects to which the nouns refer, and this will vary from context to context. Thus, in London, *cloakroomed* could not be used as in (17b), although in Washington, it could presumably be used as in (17a). It is doubtful whether most linguists would want to include information that varies from situation to situation in their descriptions of lexical items. In any case, the interpretation of *cloakroomed* appears to be a constructive process in which the listener takes the underlying noun, consults his knowledge about what its referent in this particular situation is normally used for, and then constructs an interpretation based on that normal use. So again, it appears that a process model may be necessary to account for facts that would traditionally fall within the province of linguistics.

So these are the four preliminary examples illustrating why dynamic models may be useful, perhaps mandatory, in linguistics. They demonstrate (1) how the passage of time alone interacts with judgments of grammaticality, (2) how certain linguistic rules will not be maximally general unless they treat at least some states of knowledge in the same way they treat linguistic

entities, (3) why it seems necessary for a model of language to be able to construct, consult, and alter momentary states of knowledge which in turn serve as inputs to certain linguistic rules, and (4) why perfectly ordinary interpretations of words may have to be derived from what one knows about the world, and not from a fixed lexicon. There may, of course, be ways around some of these examples, but these circumventions, we argue, may have to be taken at a loss in generality in a theory of language. To illustrate, imagine that certain constructions, because they refer to less accessible aspects of knowledge, take longer to interpret than others. Concomitantly, these same constructions will tend to be less acceptable than others because acceptability is in most cases simply a reflection of how quickly one can construct a sensible interpretation of a sentence. When no interpretation is possible, a sentence will be judged completely unacceptable. The point is that if both comprehension time and judgments of acceptability, or grammaticality, arise from the same factors, the explanation of the latter judgments, to be maximally general, should also be able to account for comprehension time. But as typically formulated, linguistic theories have no way of incorporating such notions as "instant in time" or "increment of time," to say nothing of "sequence of mental operations," as would be required to account for comprehension time. For this reason, the argument goes, the static models usually used by linguists will be inadequate and will have to incorporate certain notions of process. As an example of how this might be done we now turn to the so-called "Given-new Strategy" as a process in comprehension.

THE GIVEN-NEW STRATEGY

Fundamentally, language is a medium for the communication of information. In attempting to comprehend the utterance of another speaker, a listener has two particular goals in mind, the extraction of whatever new information the utterance may contain and the integration of this new information with information he already has. These goals are realized, we presume, by various strategies which attempt to divide the information content of a sentence into two parts: (a) given information, i.e., information already known to the listener, and (b) new information, i.e., information not already known to the listener. The structure of English, as many linguists have pointed out, actually incorporates this distinction by treating some information as presupposed and other information as asserted or implied. The listener's strategy, what we will call the Given-new Strategy, is therefore

rather simple. The listener first distinguishes between the presuppositions of the utterance and its assertion and implications. Assuming the presuppositions to convey given information, he tries to match this information with previous knowledge in memory. Then, assuming the assertion and implications to convey new information, he tries to integrate this new information at the place in memory indicated by the presuppositions. The function of given, or presupposed, information, which on the face of it seems rather unnecessary and redundant, is thus to indicate where new information is to be integrated into the memory structure. If the presuppositions of a sentence indicate several different places in memory, then the new information must make contact with all of them in order to be fully integrated and therefore fully understood. Such a process model, we will argue, is necessary if we are to account for the relative acceptability of sentences at the same time we account for the time it takes the listener to comprehend sentences in context. While a static linguistic theory may be able to account for the former, it cannot easily account for the latter.

The Given-new Strategy is based on the premise that the listener typically knows a good deal about the topic of conversation, the physical setting of the speech act, the beliefs of the speaker, etc., even before he attempts to understand a particular sentence. For lack of a more precise representation at this time, we presume that this previous knowledge is coded in memory in the form of a complex, interrelated set of primitive propositions of the type typically used by linguists in semantic representations. Thus, what the listener sees, what he thinks the speaker's beliefs are, what he has learned in the past, and what he knows from the immediately preceding sentences are all coded in such an interrelated form. Each sentence the listener then encounters contains some given information and some new. He must treat the given information as if it were an address, a pointer, or an instruction specifying where the new information is to be integrated into the previous knowledge. In brief, the Given-new Strategy is to treat the sentence as a two-part message: the given information is an address in memory, and the new information is content to be placed in the address. The strategy, therefore, consists of at least three conceptually separate stages: (1) discovering the structure of the sentence, determining what is given information and what is new; (2) finding the address in memory as determined by the given information; and (3) placing the new information at that address.

Imprecise as the Given-new Strategy is at the present time, it has several rather interesting consequences. First, it claims that the listener must attempt to match the information found in the presuppositions of a sentence with

some aspect of his previous knowledge. If the presuppositions of the sentence do not match something the listener already knows, the Given-new Strategy will fail, and he will have to try an alternative procedure. One alternative is to treat the presuppositions as new information and to set up a new structure in memory essentially unrelated to facts he already knows. Thus, if the presuppositions match previous knowledge, the sentence should be easy to comprehend; if the presuppositions do not match, or match only in some rather oblique way, the sentence should take longer to comprehend. In addition, the former sentences should be more acceptable than the latter. Second, the Given-new Strategy leads the listener to believe that the presuppositions of a sentence contain given information and the assertion new information. If, in fact, the assertion contains given information, this strategy should fail badly. The listener will be forced to alter his strategy and take more time in comprehending the sentence. In addition, he will find sentences with such inversions relatively unacceptable.

The Given-new Strategy, it should be noted, is rather different from the type of comprehension strategies proposed, for example, by Bever (1970), E. Clark (1971), and Fodor and Garrett (1967). As strategies that enable the listener to parse the surface structure more easily, these previous strategies generally refer directly to the ordering, type, and meaning of the particular formatives in surface structure and only indirectly, if at all, to the overall interpretation of the sentences. Unlike these, the Given-new Strategy assumes that the listener has already been able to parse the sentence and derive its underlying interpretation and deals instead with how the listener chooses to integrate the content of the sentence with what he already knows. In this sense the Given-new Strategy is independent of the syntactic form of the sentence. It predicts only that a given sentence will be easy or difficult to comprehend depending on what the listener already knows.

Linguistic Evidence

Before we can provide the Given-new Strategy with linguistic evidence, we need to establish the given-new structure for some representative sentences (see Firbas, 1964; Halliday, 1967, 1970; Kuno, 1972; Schachter, 1973; and others). As Halliday (1967) in particular has pointed out, the given-new structure of sentences is closely related to stress patterns, and so most printed sentences can take on a variety of given-new structures depending on how they are read. This is less true of cleft and pseudo-cleft structures as in (19) and (20), however, since their stress pattern is heavily constrained by the construction itself (barring contrastive stress):

- a. It was John who fell.
- b. Given: Someone fell. (19)
- c. New: That person was John.

- a. What John did was fall into the river.
- b. Given: John did something. (20)
- c. New: That something was: fall into the river.

The *b* and *c* portions of (19) and (20) are rough representations of the given and new components of the sentences in *a* with normal intonation. So, for example, (19) presupposes that someone fell, indicating that the speaker of the sentence presumed that the listener already knew this fact. It asserts that the someone was John, indicating that the speaker presumed that the listener did not already know that fact. (20) has a similar structure. Restrictive relative clauses, at least when not sentence final, typically carry presupposed information too:

- a. The man who fell was angry.
- b. Given: There was a man, and he fell. (21)
- c. New: He was angry.

And as Vendler (1967), for one, has noted, the definite article in certain constructions can be thought of as a type of restrictive relative clause:

- a. The man fell.
- b. Given: There is a man (under discussion). (22)
- c. New: He fell.

This contrasts with the indefinite article in (23):

- a. A man fell.
- b. Given: (Nothing) (23)
- c. New: There is a man (under discussion), and he fell.

While these examples of given-new structure are all syntactic in origin, there are other cases that arise from the properties of particular lexical items. For example, consider (24):

- a. Mary is sick too.
- b. Given: Someone besides Mary is sick. (24)
- c. New: Mary is sick.

Here the *too* brings with it a presupposition that someone else is sick—someone the listener presumably already knows about (see G. Lakoff, 1971). (Alternatively, the presupposition is that Mary is something else besides sick.) Similar examples can be found with such adverbs as *either*, *again*, *still*, *yet*, and *anymore*, e.g., *Mary isn't well either*, *Mary is sick again*, *Mary is still sick*, *Mary isn't well yet*, and *Mary isn't sick anymore*. Another type of presupposition is found in (25) (see Horn, 1969):

- a. I drink only beer.
- b. Given: I drink beer. (25)
- c. New: I don't drink anything else.

And we could give many other such examples.

As evidence for the Given-new Strategy, therefore, consider the sentence in (26):

- a. It was John who was following Mary.
- b. Given: Some person was following Mary. (26)
- c. New: That person was John.

In the Given-new Strategy, the listener is presumed to make use of this given-new structure and to assume that the speaker is essentially correct in the way he structured the sentence in (26), with one piece of information given and the other new. If the speaker has misjudged the listener, and the listener's actual knowledge contradicts that structure, comprehension should become difficult. One type of difficulty predicted by this strategy is shown in (27):

- a. Someone was following Mary. It was John who was following Mary. (27)
- b. John was following someone. It was John who was following Mary.

Note that the sequence in (27a) is relatively acceptable—discounting the rather redundant use of *Mary* in the first and second sentences—whereas the sequence in (27b) is not (assuming normal intonation). In these sequences we have set up a minimal prior context by linguistic means, supposing that the linguistic information presented before the sentence *It was John who was following Mary* will be treated by the listener as known information at the

time he attempts to comprehend this sentence. So by the beginning of the second sentence in (27a), the listener already knows that someone was following Mary. The second half of (27a) makes reference to this previous knowledge by treating it correctly as given information; it then goes on to give new information about who that someone was. In contrast, the first sentence in (27b) bears no direct relation to the given information of the second sentence nor even to the new information. The first sentence in (27b) presents the information that John was doing the following, whereas the second sentence redundantly presents the same information—at least similar information—as if it were new for the listener. At the same time the listener does not yet know that someone was following Mary, even though the second sentence in (27b) treats this information as if it were given. Thus, the unacceptability of this sequence of sentences follows directly from the Given-new Strategy, by which the listener assumes that the speaker is correct in presenting one type of information as given and the other as new.

Similar pairs of sequences—one acceptable and the other not—can be constructed for all the other given-new structures we illustrated above. Consider, for example:

- a. John did something. What he did was fall into the river.
 b. John fell somewhere. What he did was fall into the river. (28)

- a. I watched a man and a woman fall down. The man who fell was angry.
 b. I watched a man and a woman get angry. The man who fell was angry. (29)

- a. I saw a man and a woman yesterday. The man fell.
 b. I saw two people fall yesterday. The man fell. (30)

- a. John is sick. Mary is sick too.
 b. Mary is sick. Mary is sick too. (31)

In each pair of sequences, the *a* member satisfies the constraint that the given information of the second sentence is already known to the listener because of the first sentence, whereas the *b* member does not satisfy that constraint. In each case, nevertheless, the *b* sequence does present information that the listener does not already know, so the second sentence in these sequences are not entirely redundant. Indeed, the Given-new Strategy *requires* redundant

information in the second sentence. It is just that the redundant information must be found in the syntactically defined presuppositions of the sentence.

With the evidence presented so far, there are actually two possible ways of accounting for the acceptability, say, of (27a) and the unacceptability of (27b). The first is to assume, with us, that one must refer to the knowledge of the listener at the time he attempts to comprehend the sentence. The second is to assume instead that there are rules governing discourse that generate pairs of sentences, and that these rules generate (27a), but not (27b). This second explanation is quite plausible if one supposes that the grammar of English must in any case contain rules to cover the coordination of clauses with such conjunctions as *and*, *or*, and *but*, for here the problem is virtually identical to the present one (see R. Lakoff, 1971). However, such an explanation is not viable, as can be seen from the sequence in (32), which contains the previously unacceptable (27a) as a subsequence:

Mary had been followed all day, but I wasn't able to figure out who had been following her. It took me time, but I finally realized who. John was following someone. It was John who was following Mary. (32)

By the last sentence in this sequence, the listener has just the right knowledge. He knows that someone was following Mary, but he does not know who. Therefore, the information that it was John comes as something new. The problem with (27b) was not so much that the previous context indicated that John was following someone, but rather it did not clearly indicate that someone was following Mary, a presupposition needed for the optimal comprehension of the second sentence. For a similar example, consider (33) and (34).

- a. Listener: Who was following Mary?
 Speaker: It was John who was following Mary. (33)
- b. Listener: Who was John following?
 Speaker: It was John who was following Mary.

The second sequence is relatively unacceptable since the listener does not indicate that he already knows that someone was following Mary. Yet we would not want to rule out (33b) as a possible question-answer sequence, for (34), which contains (33b) as a subsequence, is perfectly acceptable, since it satisfies all the conditions about what the listener must know.

Listener: I know someone was following Mary, but who was John following? Speaker: It was John who was following Mary. (34)

In short, the acceptability of the sentence *It was John who was following Mary* is conditional on the listener having just the right knowledge, no matter how he managed to acquire it. A set of rules that blocked the sequence of sentences in (27b) or the question-answer sequence in (33b) would therefore miss the point. Such rules would fail because these same sequences are perfectly acceptable in (32) and (34).

Possible Psychological Evidence

If the Given-new Strategy led only to predictions about the relative acceptability of sequences in English, one might be able to argue for the following explanation. Each sentence requires certain conditions to hold for it to be acceptable, and some of these conditions refer to what can be presumed to be known from previous context. Sentences that violate these conditions are relatively unacceptable. Although such an explanation can account for acceptability judgments, there is clearly more to language than acceptability judgments. In particular, we would like to examine some possible temporal phenomena that appear to belong to the same family of phenomena as the acceptability judgments just examined. We will argue that for parsimony's sake there ought to be a single explanation for this whole family of phenomena, and this will require a more complex apparatus than is implied by the simple explanation of acceptability judgments given above.

The Given-new Strategy assumes that part of the time it takes people to comprehend a sentence is given over to the processes of considering the given information found in the sentence and matching that information with what is already known. Above, in presenting sentence sequences that varied in acceptability, we used extreme cases where in one sequence the given information was easy to match up with information from the previous sentence and in another it was difficult. Had we actually given such sequences to subjects and asked them to indicate the moment at which they felt they "understood" the sentences, it seems clear that the unacceptable sequences would have taken more time than the acceptable sequences. We might call this extra time a "boggle increment," since the listener would boggle at the inappropriate presentation of new information in the second sentence, and his change in strategy would take considerable time. We have not actually performed an experiment with these sequences since it seemed fairly obvious that they would show what we expect them to show.

Our interest, instead, lies in sequences with much subtler reasons for increased latencies. Consider the following sequences of sentences:

- a. John is sick. Mary is sick too.
- b. John has the flu. Mary is sick too.
- c. John has an earache. Mary is sick too. (35)
- d. John is in Agnews. Mary is sick too.
- e. John belongs to the John Birch Society. Mary is sick too.

All of these sequences are acceptable in the sense spoken of above, and none would be expected to boggle the listener. Nevertheless, we are likely to find considerable variation in how long subjects take (from the onset of the second sentence in each sequence) to claim that they understand the second sentence. The Given-new Strategy we have proposed would lead to just such variation. It is worthwhile to examine just how the strategy works on such sequences.

The sentence *Mary is sick too* has the structure shown in (36):

- a. Mary is sick too.
- b. Given: Someone other than Mary is sick. (36)
- c. New: Mary is sick.

Using the strategy proposed, the listener must attempt to match the given information against previous knowledge. This will not be equally easy for all the sequences in (35). In (35a), the match will be direct and so subjects should claim that they comprehend *Mary is sick too* quite quickly. In (35b), however, there is a slight extra step involved. The proposition *Someone else is sick* cannot be matched directly to *John has the flu* without the listener making the deduction that if John has the flu then John is sick, which matches *Someone else is sick*. This extra step should consume time, and *Mary is sick too* in (35b) should take slightly longer to comprehend than in (35a). (35c) presents a slightly different problem. The listener might not presume that if John has an earache he is sick, since although an earache is a malady, it is technically not a sickness. The listener, therefore, must go one step further and assume that the speaker (or composer of the sentence) believes that to have an earache is to be sick. Once he has made this chain of inferences the listener will be able to claim that he understands the sentence *Mary is sick too*. Thus, *Mary is sick too* should take longer to understand in (35c) than in (35a). An even longer chain of inferences seems required for (35d). As it

happens, Agnews is a town in California that quite appropriately contains a state mental hospital. The first sentence in (35d) alone indicates only that John is in the town of Agnews, but with the second sentence the listener would have to set up the following line of reasoning. The speaker must be presuming that John is sick and this would follow if John were not only in the town of Agnews but in the hospital as well. Setting up such a mental structure should take time. Consequently, *Mary is sick too* should take some time to comprehend in (35d). The final case (35e) requires a slightly different line of reasoning. In comprehending *Mary is sick too* the listener must assume that the speaker assumes that people who belong to the John Birch Society are sick, and so John, who belongs to the Society, is sick, satisfying the given information that someone else is sick. If the listener had previously assumed that the speaker himself was an ultra-conservative, then he should find it more difficult to construct this mental edifice and should therefore take even more time to claim that he understood *Mary is sick too*.

The examples in (35) show that there are many contexts in which *Mary is sick too* is a viable sentence. This sentence presupposes only that someone else is sick, which can be satisfied in a number of ways, by saying that John is sick directly or by indirect references to John's sickness. And we suppose that since to claim that he understands *Mary is sick too*, the listener must fit this presupposition into what he already knows, this sentence will require more mental effort in some contexts than others. We further assume that setting up the more complicated mental models will take more time. But is there any difference in principle between the explanation of these intermediate cases and the more extreme cases such as that in (37)?

Mary is sick. Mary is sick too. (37)

The strategy that accounts for the relative difficulty of the sequences in (35) also appears to account for the unacceptability of (37). Indeed, consider the two sequences in (38):

- a. John is well. Mary is sick too.
 b. John belongs to the John Birch Society. Mary is sick too. (38)

Sequence (38a) is obviously aberrant in isolation. The second sentence presupposes that someone else is sick, while the first asserts that John is well, and this leads to a contradiction, at least as long as the listener supposes that the first sentence was meant to satisfy the presupposition of the second. But

imagine that the listener knows that the speaker of (38b), a member of the Birch Society himself, thinks all other members of the Society are as healthy as can be. In such a circumstance the listener would judge (38b) to be just as aberrant as (38a). From the listener's point of view, these two sequences are ruled out on the same grounds. Their presuppositions contradict prior information, and according to the Given-new Strategy it does not matter whether this information is derived from previous sentences or from other sources, e.g., prior knowledge of the speaker. It is for these reasons that it seems artificial to separate the explanation of judgments of acceptability from the explanation of comprehension difficulty. In the instances we have examined, they have the same source.

Definite noun phrases will also lie on a continuum of comprehension time depending on the context they appear in. In its principle use, *the* indicates that the noun following is presupposed to exist and is in fact known to the listener. Given this rough characterization, consider the sequences in (39):

- a. John brought some apples and oranges home from the grocery store. He gave the apples to me.
- b. John brought some groceries home from the grocery store. He gave the apples to me. (39)
- c. John came home from the grocery store. He gave the apples to me.
- d. John just arrived. He gave the apples to me.

Strictly speaking, it is only (39a) that allows the listener to be absolutely certain that *the apples* refers to something John brought with him. The sequences from (39b) to (39d) become increasingly more difficult because it becomes less and less apparent what apples the speaker could be referring to. Nevertheless, in (39b) the listener would probably assume that *the apples* were part of the groceries. In (39c) he would probably assume that John had brought some apples from the grocery store. In (39d) he would be forced to construct some plausible model of the situation in which there were apples. Thus, if we asked people to indicate how long it took them to "understand" the second sentence in each of these sequences, this time should vary with the number of steps required in building a model of the situation consistent with the definite noun phrase of the second sentence.

Definite noun phrases of the type we have been illustrating refer to particular objects or events, and there are many ways, both linguistic and pragmatic, that make this function easy or difficult. Chafe (1972), for example,

has pointed out that people can assume that each room will have a ceiling and a floor, a car will have only one steering wheel, a house will have doors, and so on, and thus there is no need to posit the previous existence of such highly predictable objects. Consider the sentences in (40):

- a. Once in the room, Max stared at the ceiling in amazement.
- b. Herman got into his car and sat behind the steering wheel. (40)
- c. Jason stepped out into the light of the sun.

In these sentences the ceiling, the steering wheel, the light, and the sun can be assumed by the speaker to be known to the listener without any prior mention of these objects. Note how ridiculous it would be to have to say, "There is a sun, and it gives off light. John stepped out into the light of the sun." The examples in (39) will be more difficult to comprehend simply because the listener cannot automatically assume the existence of the *the apples* given, say, the sentence *John just arrived*.

CONCLUDING REMARKS

At this point we wish to contrast two points of view that have been implicit in this paper. The static one we will identify as the Grammarian's and the dynamic one, as the Cognitivist's. We will attempt to show how the Grammarian and the Cognitivist make quite different assumptions about the nature of language, and hence about how it should be studied.

The Grammarian's main sources of data originally consisted of judgments of grammaticality, ambiguity, and anomaly. These sources were stressed particularly strongly in Chomsky's *Syntactic Structures* and have remained foremost in the Grammarian's enterprise ever since. The grammar was supposed to generate all and only the grammatical sentences of a language and to generate two or more derivations for each ambiguous sentence. This rather narrow view of linguistics slowly gave way, within the transformational school, to the notion that the grammar should also be able to map sound onto meaning. This led to another source of data for the Grammarian, namely, semantic interpretations of sentences with known phonetic properties. Nevertheless, the notion of a grammar has changed very little. The Grammarian today sees his task as one of specifying what a person knows about his language—as reflected in judgments of grammaticality and of what a sentence means—and he does not worry about how this grammar, this theory of linguistic knowledge, might fit into a more general theory of behavior.

For the Cognitivist, the Grammarian's enterprise has seemed odd indeed. First, the Cognitivist has been puzzled by the Grammarian's choice of primary data—especially judgments of grammaticality. Although it is important at some point to be able to account for grammaticality judgments, surely they are not the *raison d'être* for the study of language. We do not speak in order to be grammatical; we speak in order to convey meaning. We do not attempt to comprehend speech in order to detect violations of grammaticality; we comprehend in order to detect meaning. In so far as comprehension and production go awry, we are willing to say that a sentence is somehow unacceptable, or ungrammatical, and these processes can go awry in a variety of ways. But surely it is more important to account for the essential function of language, namely, how it is used to convey meaning. It is apparent from recent work in linguistics that more and more linguists are moving in this direction, forsaking simple grammaticality judgments for more informative judgments of what sentences mean. Yet there is at least some of the Grammarian left in almost every linguist, for most linguists still have the lingering worry that they should be constructing a static device able to generate all and only grammatical sentences.

Second, the Cognitivist has been perplexed at how the Grammarian can believe so fervently that his theory of linguistic knowledge—the grammar—will fit nicely into a broader theory of performance. Chomsky (1965), for example, asserts that the aim of a generative grammar is “to characterize in the most neutral possible terms the knowledge of the language that provides the basis for actual use of language by a speaker-hearer.” “No doubt,” he goes on, “a reasonable model of language use will incorporate, as a basic component, the generative grammar that expresses the speaker-hearer's knowledge of the language [p. 9].” But this hope is based more on faith than on fact. The Cognitivist sees very little, if anything, in the linguistic method or its evaluation measures that would guarantee that the grammar will actually be a component of such a model of language use. The Grammarian must present evidence, not speculation, that might assure the eventual success of this approach.

At this point an analogy to chess might be useful. Imagine that Dr. Boris Fischer was a famous “cognitive psychologist,” in Chomsky's (1968) sense, and that he wanted to study how people played chess. But because he had been told that chess was a very complex game, he decided to simplify his task and examine one part of the game at a time. His decision was to study only the player moving the black pieces and to ignore the other player and the white pieces. In carrying out his research he elicited from this player all the

nonjudgmental data he could. He got the player to show him not only the move he actually would make, but also all possible moves at each point in the play. From this he was able to describe the possible moves and to specify some other general principles of the game. To his dismay, however, there was one phenomenon that remained quite mysterious. At certain points in the play certain pieces would disappear, never to appear again until the next game, and these points were quite unpredictable; the only piece that did not disappear was the King. This brought Dr. Fischer's work to a standstill. There was no coping with this mystery. Knowing what we do about chess, of course, this is hardly a mystery, and it could have been solved merely by looking at the other player and his pieces too. By his approach to the game, there was obviously only a limited amount Dr. Fischer could learn about chess beyond the elementary facts about possible moves. Because black and white are so inextricably tied to each other, and because Dr. Fischer's methods precluded his examining this interaction, there was also no guarantee that the generalizations he made were really correct. What he attributed to black and its behavior could actually be a consequence of the interaction between black and white; and such generalizations would be spurious.

The relation between grammar and language use, we suggest, may have many of the same properties as the relation between black and white in chess. Although one can discover a lot about language by studying grammar alone, it is intrinsically impossible to discover all without simultaneously looking at language use. Furthermore, many of the generalizations contained in the grammar may be spurious in that they are attributing to knowledge factors that are really a consequence of the interaction between knowledge and language use. In short, the study of language as a grammar may be inherently restrictive and misleading. Only by studying both at the same time will these errors be avoided.

As these arguments suggest, the Cognitivist views the study of language in quite a different light. He regards its central goal to be the specification of the processes by which people comprehend and produce speech, not the specification of a grammar that generates grammatical sentences or maps sound onto meaning. He regards grammaticality and sound-to-meaning mappings as derivative phenomena only, facts to be accounted for as by-products of the more central processes of comprehension and production. Consider comprehension. If the Cognitivist could specify the process by which the listener took in phonetic material, combined this with previous knowledge, knowledge of the context and rules of conversation, thereby constructing a

semantic representation for that sentence, then he would be able to account for the major results in linguistics as by-products of this process. (1) He would be able to predict grammaticality judgments from certain failures in the process of comprehension, and this is as it should be. Note that such judgments, as normally elicited, are not a product of speech production and hence have nothing *per se* to do with the process of production. (2) He would be able to predict acceptability judgments in the same way, except that these would be restricted to those failures in the comprehension process arising from memory limitations, incorrect parsings of the surface, and other "performance" factors usually ruled out as sources of grammaticality judgments. Indeed, the distinction between grammaticality and acceptability judgments is difficult to draw, if it can be drawn at all (see Watt, 1973a, 1973b); whereas this causes problems for the Grammarian, it fits neatly into the Cognitivist's conception of a single system. (3) He would be able to predict judgments of ambiguity, semantic anomaly, and contextual inappropriateness on the basis of the semantic representations produced by this process. (4) Finally, he would be able to account for the rules mapping sound onto meaning, since they would be inherent in the mental states and transitions from state to state, which are part and parcel of the comprehension process itself.

This is not to say that the Cognitivist considers comprehension to be a homogeneous process. There would undoubtedly be subclasses of rules and mental states corresponding roughly to phonology, syntax, semantics, and pragmatics. But the Cognitivist is not forced to make the strong assumptions the Grammarian has tended to make about the nature of language in relation to its use. First, the Cognitivist need not maintain the, for him, rather improbable assumption that linguistic knowledge can even theoretically be separated from the processes by which this knowledge is evinced—i.e., comprehension and production. He finds it far more likely that such knowledge is intrinsic to the actual workings of the two processes and therefore cannot be abstracted out and characterized separately. Second, even if the Grammarian were correct in assuming that linguistic knowledge is separable from language use, the Cognitivist need not maintain the even more dubious assumption that this knowledge can be *studied* independently of comprehension and production. Earlier we examined examples where primary linguistic data—grammaticality judgments—appeared to be impossible to untangle from the effects of the process by which listeners comprehend the sentences in question. This is going to be true in general. Having an adequate theory of linguistic knowledge necessitates having an adequate theory of

comprehension so that the empirical consequences of one can be separated from the empirical consequences of the other. Third, the Cognitivist need not make the rather strong a priori assumption that linguistic knowledge is different in kind from other types of knowledge—even in the trivial sense that linguistic knowledge is that which is a prerequisite for the use of language. As suggested by several examples given above, it seems indefensible in many cases to distinguish between knowledge gained from previous linguistic context (e.g., anaphors to pronouns) and knowledge gained from non-linguistic sources (e.g., deictic anaphors to pronouns). It appears that the momentary knowledge and beliefs of the speaker must be treated as mental entities of the same type as linguistic knowledge; otherwise, there could not be rules that refer to the two types of knowledge interchangeably. Fourth, the Cognitivist need not maintain, as the Grammarian does, that the temporal and sequential properties of the phonetic manifestations arise from a system that contains no mental events occurring in time. As written now, linguistic theories produce as their last stage a string of symbols in a temporal sequence without making use of any temporal devices previous to this production. The Cognitivist maintains that the temporal order of the surface of language arises from an intrinsically time-based process, one motivated by considerations of how long various mental processes take, how quickly memory fades, what mental operations are prerequisite to others, and so on.

Finally, the Cognitivist could argue that there are several positive advantages to his approach. Not only would he be able to account for grammaticality judgments, acceptability judgments, judgments of synonymy, ambiguity, and anomaly, and the interpretations people offer for sentences, but he would be able to account for how long people take in comprehending a sentence, what kind of mistakes they make in comprehension and production, how individual differences in processing capacities affect comprehension and production, and many other such phenomena. In brief, he would be able to take into account a much broader spectrum of empirical facts than is now deemed appropriate in the study of language.

We can illustrate the Cognitivist's position better perhaps with a specifically Cognitivist complaint. Linguists have typically viewed their task as one of explaining *why* people find certain sentences acceptable and others unacceptable, or *why* people interpret a sentence this way and not that. To the Cognitivist, however, one cannot give an answer to *why* without giving an answer to *how*: How—by what process—do people decide that certain sentences are unacceptable? How do people come to interpret a sentence in the way they do? If the Cognitivist knew the answer to these *how* questions,

then he would automatically know the answer to the *why* questions, but not vice versa. Since the Cognitivist is ultimately interested in explaining linguistic processes as well as linguistic knowledge, it seems more fruitful, then, to attack the broader question, letting the narrower one take care of itself.

As an illustration, consider the Complex NP Constraint proposed by Ross (1967). Consider (41):

*I know who Eva had a book that was once owned by. (41)

This sentence is ungrammatical according to the constraint given in (42):

The Complex NP Constraint: No element in a sentence dominated by a noun phrase with a lexical head noun may be moved out of that noun phrase by a transformation. (42)

Note that in (41) *who* has been moved out of the noun phrase *a book that was once owned by who*, and since the latter has a lexical head (*a book*), the constraint applies to block the sentence. Although this constraint is able to characterize *why* (41) is ruled out, and it even does so in a very general way, it clearly does not specify how people come to judge (41) as unacceptable. The Cognitivist would therefore claim that the Complex NP Constraint, as it stands now, is not sufficient. The judgment that (41) is unacceptable is itself only the end product of a complex mental process that needs to be accounted for. For the constraint to be sufficient, it would have to be shown to play a part in this process, or in the knowledge made reference to in this process, in exactly the form Ross has given it in (42). But one can imagine, for example, at least three rough ways the listener could judge (41) to be unacceptable:

(A) The listener works by analysis-by-synthesis, generating sentences by generative transformational rules until he comes up with a matching sentence; since a matching sentence will always be blocked by (42), he will never generate one, and as a consequence he will judge the sentence to be ungrammatical.

(B) The listener works from left to right in the sentence, and having found *who*, he looks for a missing case function everywhere but in NPs with lexical heads; since he cannot find one, he judges the sentence to be ungrammatical.

(C) The listener works from left to right in the sentence, and having found *by* with a missing object, he attempts to recover the missing object but

only from the complex noun phrase containing *by*; since he cannot find such an object, he judges the sentence to be ungrammatical.

Note that all three of these processes (at least when formulated precisely) are *consistent* with the Complex NP Constraint, but only Process A uses the constraint as formulated in (42). Processes B and C would lead one to reformulate the Complex NP Constraint in quite different terms—in these cases in the form of a process model specifying what the listener is actually doing in making his judgment. The point here is rather simple. While the Complex NP Constraint in (42) is consistent with the fact that (41) is judged to be ungrammatical, it may not fit into the process by which the listener comes to make this judgment, while some other formulation of the Complex NP Constraint may. The latter, obviously, ought to be preferred, but the decision of which formulation of the constraint does fit into the process will require the use and testing of process models within linguistics.

Let us consider another similar example, the Coordinate Structure Constraint also proposed by Ross (1967). Sentence (43)

*Who did Bill see Mary and? (43)

is ruled out by the constraint given in (44):

The Coordinate Structure Constraint: In a coordinate structure, no conjunct may be moved, nor may any element contained in a conjunct be moved out of that conjunct. (44)

Here again, however, the constraint, as presently formulated, is independent of the process by which the grammaticality judgment has presumably been made. And again, there are several ways the listener could judge (43) to be ungrammatical. He could do so only upon failing to find two complete conjuncts adjoined to the *and*, or only upon failing to find a missing case function for the *who* (where it is forbidden to look inside a coordinate structure), or both. Regardless of which it is, the listener will make the judgment only when there is a failure in one of the ordinary procedures he uses in the process of comprehension, and the Coordinate Structure Constraint should reflect this. Note that there are really quite low level rules the listener could use to exclude (43). It is not simply the case that conjuncts cannot be *moved* out of coordinate structures, but rather coordinate structures must always have conjuncts adjoined to both sides of the *and*. Thus (45)

*Bill saw Mary and.

(45)

can be judged as unacceptable only by inspection of the constituent *Mary and*, which is incomplete and therefore unacceptable. The listener does not need to know that the second conjunct had been moved elsewhere, as in (43), or even that there ever was a second conjunct. However this issue is to be resolved, our point is a simple one. The Coordinate Structure Constraint has not been formulated with the idea that it must eventually be referred to by, or be a part of, the process of comprehension or production, and therefore there is no evidence that it is in the correct form by this criterion. To get this evidence the linguist will have to develop or make reference to process models of language use.

And so our closing remarks come to an end. We have deliberately tried to be contentious about the goals linguists have set for themselves and about the explanations they are willing to be content with. But we see our remarks as particularly important given the recent emphasis Chomsky and other linguists have placed on linguistics as a part of "cognitive psychology" (Chomsky, 1968) and on the grammar as characterizing the competence a person "puts to use in producing and understanding speech" (Chomsky, 1970). We do not mean to disparage the linguistic enterprise and especially the tremendous number and variety of facts linguists have tried to systematize in the last decade or so. It is simply our conviction—and we have tried to motivate our convictions in this paper—that the Grammarian's point of view, like Dr. Fischer's in his study of the game of chess, will ultimately be too narrow and will lead to erroneous explanations. What is needed, we have argued, is a more perspicuous approach to language, with comprehension and production more central to the concern, and with the notion that linguistic knowledge, if it can be so characterized, is the servant, not the master, of these two processes.

REFERENCES

- Bever, T. G. The cognitive basis for linguistic structures. In Hayes, J. R. (Ed.) *Cognition and the Development of Language*. New York: Wiley, 1970.
- Bever, T. G. & Langendoen, D. T. A dynamic model of the evolution of language. *Linguistic Inquiry*, 1971, 2, 433-464.
- Chafe, W. L. Discourse structure and human knowledge. In J. B. Carroll and R. O. Freedle (Eds.) *Language Comprehension and the Acquisition of Knowledge*. Washington, D.C.: V. H. Winston, 1972.
- Chomsky, N. *Aspects of the Theory of Syntax*. Cambridge: M.I.T. Press, 1965.
- Chomsky, N. *Language and Mind*. New York: Harcourt, Brace, World, 1968.

- Chomsky, N. Remarks on nominalizations. In R. A. Jacobs and P. W. Rosenbaum (Eds.) *Readings in English Transformation Grammar*. New York: Ginn and Company, 1970.
- Clark, E. V. On the acquisition of the meaning of "before" and "after." *Journal of Verbal Learning and Verbal Behavior*, 1971, 10, 266-275.
- Clark, H. H. & Chase, W. G. On the process of comparing sentences against pictures. *Cognitive Psychology*, 1972, 3, 472-517.
- Firbas, J. From comparative word-order studies. *Brno Studies in English [Opera Universitatis Brunensis 93]*, 1964, 4, 111-126.
- Fodor, J. A. & Garrett, M. Some syntactic determinants of sentential complexity. *Perception and Psychophysics*, 1967, 2, 289-296.
- Fromkin, V. A. The non-anomalous nature of anomalous utterances. *Language*, 1971, 47, 27-52.
- Fromkin, V. A. On the reality of linguistic constructs: Evidence from speech errors, *U.C.L.A. Working Papers in Phonetics*, March 1972, 103-110.
- Garrod, S. & Trabasso, T. A dual-memory information processing interpretation of sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, 1973, 2, 155-167.
- Geis, M. & Zwicky, A. M. On invited inferences. *Linguistic Inquiry*, 1971, 2, 561-566.
- Gordon, D. & Lakoff, G. Conversational postulates. *Papers from the Seventh Regional Meeting, Chicago Linguistic Society*, 1971.
- Green, G. M. Some observations on the syntax and semantics of instrumental verbs. *Papers from the Eighth Regional Meeting, Chicago Linguistic Society*, 1972, 83-97.
- Grice, H. P. The logic of conversation. William James Lectures, Harvard University, 1967.
- Halliday, M. A. K. Notes on transitivity and theme in English: II. *Journal of Linguistics*, 1967, 3, 199-244.
- Halliday, M. A. K. Functional diversity in language as seen from a consideration of modality and mood in English. *Foundations of Language*, 1970, 6, 322-361.
- Heringer, J. T. Some grammatical correlates of felicity conditions and presuppositions. Ph.D. Dissertation, Department of Linguistics, Ohio State University, 1971.
- Horn, L. R. A presuppositional analysis of only and even. *Papers from the Fifth Regional Meeting, Chicago Linguistic Society*, 1969.
- James, D. Some aspects of syntax and semantics of interjections. *Papers from the Eighth Regional Meeting, Chicago Linguistic Society*, 1972.
- Kimball, J. P. Seven principles of surface structure parsing in natural language. *Cognition*, 1973, 2, 15-47.
- Kuno, S. Functional sentence perspective. *Linguistic Inquiry*, 1972, 3, 269-320.
- Lakoff, G. The role of deduction in grammar. In C. J. Fillmore and D. T. Langendoen (Eds.) *Studies in Linguistic Semantics*. New York: Holt, Rinehart, Winston, 1971.
- Lakoff, G. Some thoughts on transderivational constraints. In B. Kachru, R. B. Lees, Y. Malkiel, and S. Saporta (Eds.) *Papers in Linguistics in Honor of Henry and Renee Kahane*. Urbana: University of Illinois, 1973.
- Lakoff, R. If's, and's and but's about conjunction. In C. J. Fillmore and D. T. Langendoen (Eds.) *Studies in Linguistic Semantics*. New York: Holt, Rinehart, Winston, 1971.
- Lakoff, R. Language in context. *Language*, 1972, 48, 907-927.

- McCawley, J. D. Prelexical syntax. *Monograph Series in Languages and Linguistics*, 1971, 24, 35-56.
- Morgan, J. L. Syntactic form and conversational import—some interactions. In D. Cohen (Ed.) *Papers from the UWM Linguistics Group First Annual Symposium on Limiting the Domain of Linguistics*. University of Wisconsin at Milwaukee, 1972.
- Ross, J. R. Constraints on variables in syntax. Unpublished Ph.D. Dissertation, Department of Modern Languages and Linguistics, MIT, 1967.
- Sadock, J. M. Whimperatives. In J. Sadock and A. Vanek (Eds.) *Studies Presented to Robert B. Lees by his Students*. Champaign: Linguistic Research, Inc., 1970.
- Sadock, J. Speech act idioms. *Papers from the Eighth Annual Meeting, Chicago Linguistic Society*, 1972.
- Schachter, P. Focus and relativization. *Language*, 1973, 49, 19-46.
- Sternberg, S. Memory-scanning: Mental processes revealed by reaction-time experiments. *American Scientist*, 1969, 57, 421-457.
- Watt, W. C. On two hypotheses concerning psycholinguistics. In J. R. Hayes (Ed.) *Cognition and the Development of Language*. New York: Wiley, 1970.
- Watt, W. C. Competing Economy Criteria. *Social Science Working Papers 5*, University of California, Irvine, 1972. Also in *Current Problems in Psycholinguistics*. Paris: Centre National de la Recherche Scientifique, 1974.
- Watt, W. C. Late lexicalizations. In J. Hintikka, J. Moravcsik, and P. Suppes (Eds.) *Approaches to Natural Language*. Dordrecht-Holland: Reidel, 1973. (a)
- Watt, W. C. The indiscreteness with which impenetrables are penetrated. *Social Science Working Papers 20*, School of Social Sciences, University of California, Irvine, 1973. (b)