WARNING

Concerning Copyright Restrictions

The copyright law of the United States (Title 1.7, United States Code) governs the making of photocopies or other reproductions of copyrighted material.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or reproduction.

One of three specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship or research.

If electronic transmission of reserve material is used for purposes in excess of what constitutes “fair use”, that user may be liable for copyright infringement.

This policy is in effect for the following document:

Chafe, Wallace L, Discourse, consciousness, and time excerpt: / from Discourse, consciousness, and time: the flow and displacement of conscious experience in speaking and writing , SOURCE: Chicago; University of Chicago Press; 1994; pp.161-185.

Selection falls within Fair Use for 2004 Fall ONLY.

13
Some Alternative Approaches to Information Flow

The work of numerous researchers overlaps significantly with what has been set forth in chapters 5–12. These other lines of research are rich and varied, and they illustrate well one of the points made in chapter 2–that language and the mind offer a vast territory for exploration, with no easy answers. It would be impossible to review all related work here, nor am I able to do justice even to the few examples I discuss. Nevertheless, I hope this chapter will provide some helpful comparisons with certain other major contributions to the relevant "literature." I have Chosen several approaches that seem especially relevant to the present work, to the extent that the reader may be justified in wondering just what the similarities and differences are. Each has received considerable attention and has had significant influence. I hope there will be agreement that our mutual concerns involve dynamic changes in thought and language—in that sense a flow—and that these changes involve changes in the status of what may be called information, in the broad sense of negative entropy within the mind. Thus the term *information flow* in the title of this chapter seems appropriate as a way of embracing not only the consciousness-based approach of the preceding chapters, but also approaches in which consciousness has been left out of the picture.

Functional Sentence Perspective

The first scholarly tradition to bring information flow (in this broad sense) to the forefront of linguistic research, and to investigate it systematically and productively over many years, has been centered in Czechoslovakia. It has included, among other scholars, Vilém Mathesius, František Daneš, Josef Vachek, and Jan Firbas. Firbas has been and remains an especially active representative of this tradition, which has been labeled *functional sentence perspective*, a term derived from Mathesius's German term *Satzperspektive* (Mathesius 1929). An extended discussion of Firbas's approach has recently become available (Firbas 1992); there exists also a briefer and useful overview (Firbas 1986).

Functional sentence perspective has been driven by the insight that linguistic elements vary in their degree of *communicative dynamism*,

characterized as "the relative extent to which a linguistic element contributes towards the further development of the communication" (Firbas 1992, p. 8). "It is an inherent quality of communication and manifests itself in constant development towards the attainment of a communicative goal; in other words, towards the fulfillment of a communicative purpose. Participating in this development, a linguistic element assumes some position in it and in accordance with this position displays a degree of communicative dynamism" (p. 7).

Communicative dynamism thus assumes, but does not explicitly develop, a theory of language use–one in which speakers, when they say something, have a communicative purpose, with the elements of their language contributing to that purpose to a greater or lesser degree. Discussions of functional sentence perspective rely on introspections regarding the goals of communicative acts, but they avoid any broader social or cognitive commitments: "I have not studied the relationship between degrees of CD and their counterparts in the mind of the language user, but I do not think that the language user is unaware of the development of the communication" (Firbas 1992, p. 107). We are left to guess what is meant by communicative purpose from constructed exchanges such as the following (Firbas 1986, p. 42):

(1) a What about Peter?
b He lias flown to Paris.

The purpose of the response in (1)b is said to be "to state the destination of Peter's flight." One can easily imagine other purposes this imaginary speaker might have had, but the lack of context leaves the question open. In any case, the word *he* is said to contribute the lowest degree of communicative dynamism, *has flown* an intermediate degree, and *to Paris* the highest degree, since it directly expresses Peter's destination. One of the findings of functional sentence perspective has been that, all other things being equal, the order of words in a sentence corresponds to an increase in communicative dynamism. To that extent, then, functional sentence perspective provides a functional explanation for word order.

But that is far from the whole story. Communicative dynamism is said to be "determined by" four factors identified as (a) linear modification, (b) the contextual factor, (c) the semantic factor, and (d) prosodic prominence (in spoken language only). Linear modification is a term taken from Dwight Bolinger (1952, p. 1125; also Bolinger 1965, p. 288): "gradation of position creates gradation of meaning when there are no interfering factors." Although Bolinger used this principle to explain a somewhat different phenomenon, Firbas has used it to capture the relation between word order and communicative dynamism, as illustrated in (1)b. The contextual factor involves " retrievability/irretrievability from the immediately relevant context" (Firbas 1992, p. 21), thereby creating the opposition *context-dependent* versus *context-independent.* It is evidently in part a matter of identifiability, but more a matter of activation cost. The semantic factor involves what are called *dynamic functions,* elements in a rather complex theory that attributes a semantic basis to parts of speech, grammatical relations, and other grammatical phenomena, an effort with which I am in principle much in sympathy. The well-known distinction between *theme* and *rheme* is included here. Finally, the study of prosodic prominence integrates functional sentence perspective with British intonation studies.

In terms of the present work, functional sentence perspective is a mixture of several things. If it were translated to accord with this work, it would say that there is a single dimension (communicative dynamism) that is "determined by" a complex interaction of word order (linear modification), activation cost and identifiability (the contextual factor), various semantic elements and relations that underlie grammar (the semantic factor), and prosody. Communicative dynamism probably corresponds most directly to a blend of what I have been calling referential importance (chap. 7) and newsworthiness (chap. 12). It is thus on the same plane as, but distinct from, activation cost and identifiability (chaps. 6 and 8). Word order and prosody, on the other hand, are aspects of linguistic expression. Semantics involves still other aspects of thought and language. Viewed in these terms, communicative dynamism is not a unified phenomenon.

Of particular interest to the present work is Firbas's recognition of the special status of (a subset of) what I have been calling low-content verbs—"verbs or verbal phrases that explicitly convey the meaning of appearance or existence on the scene" (Firbas 1992, p. 60). In constructed examples such as:

(2) A boy came into the room.

"the subject is context-independent and conveys the information towards which the communication is perspectived" (that is, it exhibits the highest degree of communicative dynamism). "The notional component of the verb introduces this information into the communication and in this respect recedes into the background" (p. 59). I would assign such verbs to the presentative subset of low-content verbs. It is interesting to see that quite different motivations led both avenues of research to assign a special place to verbs of this kind. Functional sentence perspective was motivated by the desire to assign a lower degree of communicative dynamism to verbs whose subjects carry a higher degree, as in (2). The recognition of low-content verbs in chapter 9 emerged from an examination of potential counterexamples to the one new idea hypothesis, some of which exhibit these verbs.

Viewed from the perspective of this work, then, functional sentence perspective unites several distinct discourse functions (activation cost, identifiability, referential importance, newsworthiness) within a single dimension of communicative dynamism. It stops short of understanding these matters within a larger socio-cognitive frame of reference and ignores the role of consciousness. More positively, it has pioneered in examining some of the basic questions in this area and has brought a variety of provocative issues to the forefront of research.

Functional Grammar

Michael Halliday, who has presented his carefully developed ways of understanding these aspects of language under the label *functional grammar*, has long been concerned with many of the aspects of language that are treated to this book. His work is highly ramified and covers far too many aspects of language to be summarized here. I will limit the discussion to just a few areas that are especially clearly presented in Halliday (1985b), on which most of the following remarks are based (see also Halliday 1985a). A recent sympathetic discussion of relevant aspects of his approach has appeared in Vande Kopple (1991).

Halliday has been one of the few linguists who have for some time been fully aware that conversational language and written language have different properties (e.g., Halliday 1987), and he has stressed that there is much to be gained from observing natural spoken language: "Perhaps the greatest single event in the history of linguistics was the invention of the tape recorder, which for the first time has captured natural conversation and made it accessible to systematic study," for "it is in spontaneous, operational speech that the grammatical system of a language is most fully exploited" (Halliday 1985b, pp. xxiii–xxiv).

Like the present work, Halliday has searched for correspondences between linguistic elements and their functions. One such element is the *tone group*. Importantly, "the tone group . . . is not only a phonological constituent; it also functions as the realization of something else, namely a quantum or unit of information in the discourse. Spoken discourse takes the form of a sequence of *information units*. . . . The information unit is what its name implies: a unit of information. Information, as this term is being used here, is a process of interaction between what is already known or predictable and what is new or unpredictable" (Halliday 1985b, pp. 274–75).The "already known or predictable" is what Halliday calls *given*, as opposed to the "unpredictable" or *new*. He amplifies these characterizations by explaining that "the significant variable is; information that is presented by the speaker as recoverable (Given) or not recoverable (New) to the listener" (p. 277). Further, "the meaning [of given] is: this is not news." "The meaning [of new] is: attend to this; this is news."

The similarities and differences should be evident. Both the present work and Halliday's recognize the fundamental importance of what I have been calling intonation units, Halliday's tone groups. Both recognize that these units include some elements that are in some sense given and others that are in some sense new. But there are differences in what the terms *given* and *new* are taken to mean. The present work understands these terms with relation to the speaker's assessment of activation cost in the mind of the listener. Halliday also recognizes that "Given + New is listener-oriented" (1985b, p. 278), but he characterizes these properties in terms of recoverability. "What is treated as recoverable may be so because it has been mentioned before; but that is not the only possibility, It may be something that is in the situation, like *I* and *you*; or in the air, so to speak; or something that is not around at all but that the speaker wants to represent as Given for rhetorical purposes" (p. 277). Such a statement approaches but does not coincide with an explanation in terms of presence in active consciousness. Halliday does not recognize a degree of activation cost (or recoverability) that is intermediate between given and new, discussed here in terms of accessibility. He comes close to recognizing the one new idea constraint–"an information unit consists of an obligatory New element plus an optional Given" (p. 275), though it is not stated as such. The greatest divergence from the present work, however, appears in his treatment of subjects and themes.

To understand his use of these two terms, it is necessary to recognize the importance to all of Halliday's work of positing three "kinds of meaning," or "metafunctions," which he labels *ideational*, *interpersonal*, and *textual:* "Ideational meaning is the representation of experience: our experience of the world that lies about us, and also inside us, the world of our imagination. It is meaning in the sense of 'content.' . . . Interpersonal meaning is meaning as a form of action: the speaker or writer doing something to the listener or reader by means of language. , . . Textual meaning is relevance to the context: both the preceding (and following) text, and the context of situation" (Halliday 1985b, p. 53). Halliday sees a clause as functioning simultaneously as a *message* (the ideational function), an *exchange* (the interpersonal function), and a *representation* (the textual function). In a discussion of the following constructed sentence (p. 32),

(3) The duke gave my aunt this teapot.

the idea of the duke is said to function simultaneously as theme, subject, and *actor*. These three "functional concepts" are interpreted as corresponding to the three different modes of meaning:

(a) "The Theme is a function in the *clause* as *a message.* It is what the message is concerned with: the point of departure for what the speaker is going to say.
(b) "The Subject is a function in the *clause as an exchange.* It is the element that is held responsible: in which is vested the success of the clause in whatever is its particular speech function.
(c) "The Actor is a function in the *clause* GIS *a representation* (of a process). It is the active participant in the process: the one that does the deed." (P. 36–7)

It is by no means necessary that the same element (like *the duke* in *(3))* be simultaneously theme, subject, and actor. In *(4) this teapot* is said to be the theme, *my aunt* the subject, and *the duke* the actor:

(4) This teapot my aunt was given by the duke.

There is no need to dwell on the *actor* function. It evidently corresponds to the agent role as discussed in chapter 12, one of various semantic roles a referent can have in an event. It may be a core role, as in (3), or it may be expressed by a prepositional phrase, as in (4). I am uncertain why the semantic role of actor should be associated with the textual function of language. However that may be, it is the subject and theme functions that contrast most noticeably with the interpretations set forth in this book.

A subject is said to be

> something by reference to which the proposition can be affirmed or denied. For example, in *the duke has given away that teapot, hasn't he?* . . . the Subject *the duke* specifies the entity in respect of which the assertion is claimed to have validity. It is the duke, in other words, in whom is vested the success or failure of the proposition. He is the one that is, so to speak, being held responsible — responsible for the functioning of the clause as an interactive event. The speaker rests his case on the *duke + has,* and this is what the listener is called on to acknowledge. (P. 76)

In attributing subjecthood to "the clause as an exchange," Halliday sees it as "setting something up so that it can be caught, returned, smashed, lobbed back etc." (p. 76n.). For example, listeners might respond to (3) by saying *No he didn't,* thus showing that for them the proposition failed.

The tennis ball metaphor is related to Halliday's prescription for identifying a subject: "The Subject, in a declarative clause, is that element which is picked up by the pronoun in the tag" (p. 73). The fact that *he* in the tag at the end of (5) refers to the duke provides a simple way of identifying *the duke* as the subject of what precedes:

(5) The duke gave my aunt this teapot, didn't he.

The fact that *Yes he did* or *No he didn't* are so closely related to the tag is taken as evidence for the subject's role as the expression of "something by reference to which the proposition can be affirmed or denied." This view of the function of subjects is an interesting one, but I believe it can be seen as a consequence of their role as starting points. If that is correct, it would appear that Halliday has been prevented from acknowledging subjects as grammaticized starting points because that role has been preempted by what he calls themes.

What, then, is a theme? "In English, as in many other languages, the clause is organized as a message by having a special status assigned to one part of it. One element in the clause is enunciated as the theme; this then combines with the remainder *so* that the two parts together constitute a message" (p. 38). Particularly interesting is the statement, "The Theme is the element which serves as the point of departure of the message; it is that with which the clause is concerned." This might suggest that Halliday's theme is equivalent to what in this book is called a starting point, but that is not the case: "In speaking or writing English we signal that an item has thematic status by putting it first" (p. 38). Halliday says that in a sentence like (6),

(6) This teapot my aunt was given by the duke.

this teapot is the theme, while *my aunt* is the subject and thus in my terms the starting point. The theme need not be a referent at all. In the following examples the italicized initial phrases are all said to be themes (p. 39):

(7) *Once* I was a real turtle.
(8) *Very carefully* she put him back on his feet again.
(9) On *Friday night* I go backwards to bed.

Indeed, a sentence may have multiple themes, each of which may contribute either a textual, interpersonal, or ideational function. In the following example, *on the other hand* is said to be a textual theme, *maybe* an interpersonal theme, and *on a weekday* an ideational theme (p. 55):

(10) On the other hand maybe on a weekday it would be less crowded.

Halliday, then, interprets the first element in a clause as having a special functional status, labeled theme, though he allows for a sequence of themes of the type just illustrated, The function of a theme is to express what he has characterized as the starting point of a message. There is a clash of introspections here. Halliday sees starting points as expressed in the first element of a clause; I see them as expressed in subjects. Are both interpretations circular, since one says we know something to be a starting point because it occurs first and it occurs first because it is a starting

point, while the other says we know something to be a starting point because it is a subject and it is a subject because it is a starting point? We know that introspections alone can lead to different conclusions, a fact well illustrated here, but that is no reason to discard them as having no validity. I suggested in chapter 2 that language is uniquely valuable for the study of the mind because it provides a wealth of complex phenomena that can be paired with introspections. The stronger the linguistic side of the pairing, the more validity can be attached to the introspection. There is something intuitively valid in the notion of starting point, but a full understanding of the nature and role of starting points depends on the richness and relevance of their linguistic correlates.

The question involves the extent to which the starting point function is convincingly paired with initial position versus the extent to which it is convincingly paired with subjecthood. I have tried to show not only that starting points are paired with subjecthood as a grammatical status, but that starting points also conform, in conversational language at least, to discourse properties one might expect of referents functioning in that way. With respect to activation cost, they exhibit the property discussed in chapter 7 as "lightness": most are given, some are accessible, and a small residue is new but of trivial importance. Almost all subjects show identifiability. Subjects tend to be the referents from whose point of view something is expressed. They also tend to exhibit the semantic property of humanness and to perform the semantic role of agent. All of these are properties we would expect starting points to have, and thus they provide multifaceted support for the introspection.

The property of being the first element in a clause is less coherent. Such an element may be, and often is, the subject, but it may alternatively be an orientation of some kind—spatial, temporal, epistemological, textual–or sometimes a referent that is being contrasted with some other referent. To say that *on the other hand maybe on a weekday* is the starting point in

(11) On the other hand maybe on a weekday it would be less crowded.

whereas *it* is the starting point in

(12) It would be less crowded.

misses, I believe, the function of it in both sentences, and confuses the starting point function with the orienting one. It may be that newsworthiness (chap. 12), contrastiveness (chap. 6), "topichood" in the sense mentioned at the end of chapter 8, and perhaps other factors lead to the placement of an element in initial position, but only the study of natural examples in context, with their prosody, can sort these matters out.

Halliday's work has covered much the same range of phenomena as the present work. One of its most useful features has been its recognition of the importance of prosody, and especially the importance of the tone group or intonation unit. I suggest that it has arrived at different conclusions partly because of the mixed quality of its data, partly because it has not recognized the role of consciousness, and partly because it has been committed to a unitary functional role for the diverse elements that may appear first in a clause.

The Given-New Contract

Well-known and influential work of a very different sort was reported by Herbert Clark and Susan Haviland in the 1970s (Haviland and Clark 1974, Clark and Haviland 1977, Clark 1977; see also Clark and Clark 1977, pp. 95–98) and has continued to influence research on the given-new distinction. Their underlying conception was that of a "given-new contract" agreed to by the speaker and listener, a contract that was seen as one aspect of the "cooperative principle" popularized by Paul Grice (1975). One of the attractive aspects of this view was the recognition that "the speaker tries, to the best of his ability, to make the structure of his utterance congruent with his knowledge of the listener's mental world" (Clark and Haviland 1977, p. 4).

Clark and Haviland were concerned not only with the speaker but also with the listener, from whose point of view "the given-new strategy is a three-step procedure for relating the current sentence to this knowledge base. At Step 1, the listener isolates the given and the new information in the current sentence. At Step 2, he searches memory for a direct antecedent, a structure containing propositions that match the given information precisely. Finally, at Step 3 the listener integrates the new information into the memory structure by attaching it to the antecedent found in Step 2" (p. 5). These three steps were illustrated with the following unusual constructed sequence (pp. 4–6):

(13) a Someone piqued the professor.
b It was Percival who piqued the professor.

Having heard (13)a, the person who heard (13)b would begin processing it by dividing it into its given and new parts, the given being X *piqued the professor* and the new being *X = Percival.* (It was assumed that the nature of cleft sentences such as (13)b was to distribute given and new information in this way.) Second, the listener would search his or her memory for a unique antecedent that matched the given information, finding it in what had been acquired from the previously heard sentence (13)a. Third, the listener would integrate the new information in (13)b

with this given information by replacing *X* with *Percival,* thus now knowing that *someone = Percival.*

Sometimes, however, the listener would not be able to find a direct match for the given information within knowledge already possessed, but would be forced to construct a *bridge* between what was known and what was treated as given:

(14) a Ed was given lots of things for his birthday.
b The alligator was his favorite present.

The alligator in (14)b was said to express given information that had no direct representation in the listener's knowledge structure (Haviland and Clark 1974, p. 514): "With no direct Antecedent for the Given information in the target sentence [14b], the connection between the two sentences requires an extra inferential step, something like, 'Ah, one of those "things" must have been an alligator.' " Clark and Haviland hypothesized that this bridging operation would require a certain amount of extra time, over and above whatever time would have been involved in just searching memory for a direct match for the given information. Several experiments to measure reaction times were performed to see whether they would confirm this hypothesis.

in one experiment, subjects first saw on a tachistoscope a *context* sentence like (15)a:

(15) a We got some beer out of the trunk.
b The beer was warm.

When they had read it, they pressed a black button, (15)a disappeared, and they saw a *target* sentence like (15)b. They were instructed to press a red button as soon as they understood what (15)b meant. it took them a mean time of 835 milliseconds to do that. Other subjects, instead of seeing a context sentence like the one in (15)a, saw a sequence like the following

(16) a We checked the picnic supplies.
b The beer was warm.

It took these subjects longer (1,016 milliseconds) to press the red button.

This observation was interpreted as support for the bridging hypothesis. Although the subjects *read* these sentences, Clark and Haviland *(1977,* p. 21) referred to them as "listeners." When subjects saw (16)a followed by (16)b, Clark and Haviland said, "there is no direct antecedent in the context sentence, and so the listener must build a bridge. He must draw the implicature that the picnic supplies contain a quantity of beer, and it

is that quantity that is being referred to by the given information of the target sentence. Since drawing this implicature presumably takes time, the listener should take longer to comprehend the target sentence *The beer was warm"* in the *indirect antecedent* sequence–(16)a followed by (16)b—than in the *direct antecedent* sequence—(15)a followed by (15)b.

Experiments always leave room for alternative interpretations, and in this case Clark and Haviland noticed that context sentence (16)a did not contain the word *beer,* whereas context sentence (15)a did. "The direct antecedent sequences may have been easier simply because of the repetition of the word *beer,* perhaps making the second instance of *beer* easier to comprehend" (Clark and Haviland 1977, p. 22). To see whether it was just the repetition of the word that made the difference or whether it really was the process of bridging, sequences like the following were substituted for those in (16):

(17) a Andrew was especially fond of beer.
b The beer was warm.

"Again as predicted, comprehension time for target sentences was faster for the Direct Antecedent pairs than for the Indirect Antecedent pairs, 1031 to 1168 msec. . . . These results, therefore argue that mere repetition of the critical noun is not enough to account for the results of Experiment I" (Haviland and Clark *1974,* p. 516).

Viewed from the perspective of this book, Clark and Haviland's experiments raise some interesting questions. Let us at first assume that what is involved here is what I have called activation cost, although we will shortly see reason to doubt that assumption. The sequence in (15) then illustrates a straightforward case of givenness: the referent was activated in (15)a and retained its active status in (15)b. Of course, it would have been more natural in that case for (15)b to have contained a weakly accented pronoun:

(18) a We gat some béer out of the trunk.
b Itwaswárm.

But we can accept the full noun phrase in (15)b as a not very disturbing manifestation of the psychologist's license to sacrifice naturalness for control. We might at least suppose that the subjects' auditory imagery of (15)b assigned a weak accent to the word *beer:*

(19) The beer was wárm.

In (16)b, on the other hand, the word *beer* would undoubtedly have been assigned a primary accent if it had been spoken, and must have been

imaged auditorily as having such an accent when the subjects read it silently:

(20) a We chécked the picnic supplies.
b The béer was wárm.

Thus, although they looked identical on the tachistoscope, (15)b and (16)b would have been perceived as prosodically different.

Why did it take the subjects longer to process (16)b? If the explanation is limited to activation cost, we can conclude that processing an already active referent takes less time than activating a referent that was previously in a less than completely active state. The question then arises as to whether the idea of the beer in (16)b, since it was not given, was new or accessible. A new referent would have violated the light subject constraint, so it is worth considering why the idea of the beer would have been semiactive at this point. Obviously its accessibility must have arisen through association with the idea of the picnic supplies. It is thus possible that the extra time taken to process (16)b was occupied in activating a referent that was previously semiactive and not fully active as in (15).

However, there is another and probably better way of explaining the longer reaction time. Not only was *the beer* in (16)b treated as accessible, it was also treated, through the use of the definite article, as identifiable. The contrast between (15)b and (16)b is precisely the contrast between indirect and direct sharedness as discussed in chapter 8. The idea of the beer in (15)b had already been established as a directly shared referent in (15)a. The idea of the beer in (16)b was identifiable because of the knowledge that picnic supplies are likely to include beer, The extra time it took to process (16)b would then have resulted from the reader's need to establish identifiability on the basis of indirect sharedness. In brief, Clark and Haviland's experiment might be interpreted as showing that an accessible referent takes longer to process than a given one. Alternatively, it could be interpreted as showing that an indirectly identifiable referent takes longer to process than a directly identifiable one. This second interpretation seems more likely to be correct, but it has nothing to do with the given-new distinction.

It is interesting also to give some thought to the second of Clark and Haviland's experiments, in which they found a longer processing time for (21)b as compared with (15)b. I have added accents that reflect the subjects' most likely auditory imagery:

(21) a Àndrew was especially fònd of béer.
b The beer was wárm.

Here the context sentence, (21)a, is generic and establishes only the generic idea of beer, not the idea of any particular beer. As noted in Chafe

(1974, pp. 125–27), a generic referent can establish *givenness* for any instance of the category in question. Hence, if (21)a had been followed by a sentence like

(22) He brought some beer with him.

the phrase *some beer* would have had a given referent and would thus have been pronounced with a weak accent. For the same reason, the experimental sentence (21)b would have been most naturally pronounced with a weak accent on *the beer*. But (21)b would be an odd thing to say in the context of (21)a. What is odd is the fact that the idea of the beer was treated as identifiable. The context provided by (21)a did not establish the idea of any particular beer, but only of generic beer; hence there was no particular idea to be shared. It is, in fact, interesting to observe that indirect sharedness cannot be derived from a generic referent in the same way it is derived from a particular one like *the picnic supplies*. To repeat these two examples, (24) is peculiar in a way that (23) is not:

(23) a We chécked the picnic supplies.
b The béer was wárm.

(24) a Àndrew was especially fond of béer.
b The beer was wárm.

One could imagine (23) actually occurring. The natural occurrence of (24) is doubtful, unless the particular beer had been introduced earlier.

Judging from these examples, the findings concerning (23) and (24) might be reinterpreted as follows. Sentence (23)b required extra processing time because the identifiability of the beer had to be established on the basis of indirect sharedness. Sentence (24)b required extra processing time because of a more daunting problem with identifiability — the fact that there was no basis even for *inferring* a shared referent. The reaction times did not distinguish these two quite different processes, but there may have been a ceiling on how long subjects would take to push the red button no matter what they saw. Participants in an experiment learn not to be startled by unusual language. Experiments can be helpful, but they can leave basic questions unresolved so long as they are isolated from observations of natural language, and from crucial introspective evidence as well.

I should add that Clark's contributions to discourse understanding hardly ended with the research just described, which I have discussed at length because of its direct relevance to this book and the fact that it is still frequently cited. More recently, among other lines of research, he has added to our understanding of identifiability and has been especially concerned with the collaborative nature of mutual understanding (e.g.,

Clark and Wilkes-Gibbs 1986, Schober and Clark 1989, and other papers reprinted in Clark 1992).

Hierarchies of Assumed Familiarity, Accessibility, or Givenness

Ellen Prince's 1981 article on the given-new distinction provided a taxonomy of given and new information that has been used by a number of workers in this area. It should, therefore, be of special interest to compare her way of classifying these phenomena with the way they have been treated here. Prince's taxonomy brought together within a single category several discourse properties I have treated as distinct, uniting them under the heading *assumedfamiliarity* Her exposition was based on an analogy to recipes, which may be verbalized in different ways depending on the writer's "assumptions about what the reader knows about ingredients, processes, and equipment, about what equipment the reader has available, and about what staples the reader keeps on the shelf" (Prince 1981b, pp. 234–35).

Accepting the cooking metaphor as a way of understanding assumed familiarity, we can consider first what is meant by *new:* "When a speaker first introduces an entity into the discourse, that is, tells the hearer to 'put it on the counter,' we may say that it is *new"* (Prince 1981b, I 235). (Her *entity* is equivalent to what I have been calling a *refèrent.)* Ultimately, of course, it is necessary to get behind the cooking metaphor to arrive at an appreciation of the mental states and processes the metaphor is designed to help us understand. I believe there is no good way to understand what "putting a referent on the counter" means except as a way of visualizing what happens when someone places in active consciousness a referent that was previously inactive. In other words, Prince's explanation of new information need not conflict with an explanation in terms of consciousness, which it avoids by inviting us to compare mental processing with cooking.

Prince went on to distinguish two kinds of *new* referents: "In one case, the hearer may have had to *create* a new entity, akin to going out and buying a suckling pig, in which it is *brand-new*. In the other case, the hearer may be assumed to have a corresponding entity in his/her own model and simply has to place it in (or copy it into) the discourse-model, akin to taking some staple off the shelf when its presence is suddenly taken for granted in a recipe (e.g. salt). Call this type *unused"* (1981b, p. *235)*. This distinction is almost identical to that which I have labeled *ushared* versus *shared* (chap. 8). As I have presented it, however, sharedness is independent of activation cost, to which it is related solely through the logical necessity that only new ideas can be unshared.

Whether or not a referent is assumed to be newly activated in the listener's consciousness is a different question from whether or not it is assumed to be already part of the listener's knowledge. Activation cost is manifested linguistically in such phenomena as the use of a pronoun or a full noun phrase, as well as in weak or strong prosody. Sharedness, on the other hand, is one of the components of identifiability, which is manifested in various ways, but most conspicuously in the use of the definite article. There is nothing wrong, then, with saying that new referents may be either brand-new or unused (unshared or shared), so long as the latter distinction is understood to be on a different cognitive and linguistic plane from that which defines them as new. Since referents that are *not* new are necessarily shared, the term *unused* would seem to be an appropriate way of designating only those shared referents that are new, rather than all shared referents. As a term, therefore, *unused* has the disadvantage of conflating the separate domains of sharedness and activation cost.

Prince further distinguished brand-new entities that are *anchored* from those that are not. "Brand-new entities themselves seem to be of two types: *anchored and unanchored.* A discourse entity is Anchored if the NP representing it is *linked* by means of another NP, or 'Anchor,' properly contained in it, to some other discourse entity" (1981b, p. 236). Although Prince's exposition made use of constructed sentences, she applied her taxonomy to a real conversational narrative taken from Nessa Wolfson (1982, pp. 94–95).This narrative contained several examples of anchored new entities, one of them at the very beginning:

(25) Well, I have a friend of mine called me:

The new referent expressed as *a friend of mine* was anchored because the idea of the friend was linked to the idea of the speaker with the phrase *of mine.* The anchor in such a case is usually, if not always, something other than brand-new: "In the data, all Anchored entities contain at least one Anchor that is not itself Brand-new" (Prince 1981b, p. 236).

Like sharedness, anchoring is distinct from activation cost as such. It does, however, raise some interesting questions and suggest the need for further research. It is instructive to look at the following sequence, which was part of the exchange discussed in chapter 10:

(26) a(A) ... Hòw're you dòin' with the hóuse.
b(B) ... Òh got it àll uh ... primed just abòut,
c(B) ... except twó sides of it.

The referents verbalized as *the home* in (26)a and *two sides of it* in (26)c are both new, and both can be interpreted as shared (unused). The first is unanchored, whereas the second is anchored with the phrase *of it.* We can *see* from this example that anchoring is not restricted to unshared

(brand-new) referents. But there is a larger question here that involves the manner in which new referents — whether they are shared or unshared, anchored or unanchored–are introduced into a discourse.

Speakers tend not to introduce new ideas out of the blue, but fit them in some way into the ongoing interaction. To be sure, (26)a introduced a new topic into the conversation, but from Speaker B's response it is evident that knowledge of the painting of the house was already shared. Furthermore, although *the house* was, strictly speaking, unanchored, it was linked to its context with the word *you*, which made saying *your house* unnecessary. Thus, both *the house* and *two sides of it* were linked to shared knowledge, though in different ways. The former indicated that fact overtly only through the use of the definite article, whereas the latter included an overt anchor of the sort described by Prince. When Speaker B later introduced a new referent as

(27) (B) .. Òh = that gútless ... spráyer it

the idea of this particular sprayer may have been unshared. However, the immediately preceding intonation unit took the form of Speaker A's implied question:

(28) (A) .. I thought you were gonna spráy it.

In that context, mention of a new and unshared sprayer was quite natural. The point I am making is that anchoring must be one aspect of a larger strategy–the manner in which new ideas are related to their contexts. For *any* new idea we can ask how it is linked to the context in which it is introduced. Explicit anchoring, when it is present, provides an overt indication of what the link is, but there are other kinds of links that need to be investigated too.

Although Prince did not make a point of it, the concept of anchoring has a particular relevance to identifiability. In chapter 8 I discussed various types of what I called sufficiently identifying language–language sufficient to make shared referents identifiable. One type was the creation of identifiability through modification of a category: the use of possessors, attributive adjectives, prepositional phrases, and relative clauses. For example, in Prince's illustrative narrative the speaker said:

(29) Well, try the kitchen window,

in a context where presumably *the window* alone would not have been sufficient for identifiability. Prince's examples of anchoring, in contrast, involve *nonidentifiable* referents, like *a friend of mine* in (25). It appears, then, that anchoring performs two very different functions: relating a new idea to its context and creating identifiability.

As another category in her taxonomy, Prince used the term *evoked* as

an equivalent for *given*. She explained the evoked status again with the culinary analogy: "Now, if some NP is uttered whose entity is already in the discourse-model, or 'on the counter,' it represents an *evoked* entity" (1981b, p. 236). She then distinguished *textually* evoked entities from those that are *situationally* evoked. The distinction is a matter of how givenness is established, and parallels the discussion in chapter 6 of activation through the discourse or through salient presence in the extralinguistic environment.

In addition to entities that are new and those that are evoked (or given), Prince recognized a third category: those that are *inferrable* "A discourse entity is Inferrable if the speaker assumes the hearer can infer it, via logical–or, more commonly, plausible — reasoning, from discourse entities already Evoked or from other Inferrables" (p. 236). Her example was *the driver* in

(30) I got on a bus yesterday and the driver was drunk.

As she explained, "*the driver* is Inferrable from a *bus*, plus assumed knowledge about buses, that is, *Buses have drivers.*" This type of explanation might be more appropriate to explaining the *identifiability* rather than the givenness of a referent like *the driver*, as in the case of Clark and Haviland's *the beer*. In terms of the discussion in chapter 8, it would be said that the idea of the driver is indirectly shared as a result of association with the idea of the bus. However, it would appear that Prince regarded the referent of *the driver* as having a different status with respect to "assumed familiarity" and that she was not concerned with its identifiability. She provided no culinary analogy, and as a result the relation of inferability to the rest of her taxonomy remains uncertain.

Prince went on to suggest a "preferred hierarchy or scale for what type of entity is used" (1981b, p. 245). This scale was discussed with reference to the following constructed examples:

(31) a I bought a Toyota.
b Ellen bought a Toyota.
c One of the people that work at Penn bought a Toyota.
d A person that works at Penn bought a Toyota.
e A person bought a Toyota.

Prince continued: "It seems that, if a speaker is in a position to say one of these on basis of his/her hypothesis about what the hearer knows and chooses instead to say one lower on the scale (to refer to the same individual), s/he will be seen, if found out, to have been deviant in some way (e.g. evasive, childish, building suspense as in a mystery novel). Put differently, we may say that the use of an NP representing a certain point

on the scale implicates that the speaker could not have felicitously referred to the same entity by another NP higher on the scale" (p. 245).

We can note that several disparate factors are at work in (31). For example, although the referents of *I* in (31)a and *Ellen* in (31)b might be given, the appropriateness of using one or the other would depend on who was speaking, Ellen or some third party. It is more relevant to the present discussion to leave (31)a out of account and try to imagine situations in which the subjects of the remaining sentences conveyed new information. Prince's point was that, for example, if the referent was shared (unused), (31)b would take precedence over the choices below it. That is certainly true, but it requires considerable imagination to think of situations where any of the others would be used at all. ßeyond that, to say that if the referent is unshared (brand-new) but anchored, and thus (31)d would be used in preference to (31)e, is to say nothing at all, for it is precisely the presence of an anchor that differentiates (31)d from (31)e. The hierarchy as presented combines several distinct dimensions into one, but to point that out is not to deny the value of looking at how referents are most likely to be categorized, a process in which activation cost, sharedness, familiarity, context, and other factors play a role.

Finally, it is worth noting that Prince's examination of conversational language found that "nearly all of the subjects are Evoked"–that is, given (1981b, p. 242), and that in "informal conversational discourse" the tendency is "to reserve subject position for NPs at the higher end of the scale" in (31)–that is, identifiable (p. 246). These findings, of course, help support the validity of the light subject constraint. Prince also investigated the somewhat different patterns of information flow that are observable in written language, a topic to which we will return in Chapter 22.

More recently, two other hierarchies have been suggested, bearing some resemblance to that exemplified in (31), but each different in its own way. One is the "accessibility" hierarchy set forth by Mira Ariel (1988, 1990, 1991). In terms of the present work, this kind of accessibility involves what I discussed in chapter 8 as the use of sufficiently identifying language when a speaker is verbalizing a shared referent. Ariel lists the following types of linguistic expressions, ranging from those used when identifiability calls for a more informative verbalization–a situation she terms low accessibility–to those used when a minimum amount of verbal material is sufficient–a situation termed high accessibility (Ariel 1991, p. 449):

Full name + Modifier
Full name
Long definite description
Short definite description
Last name

First name
Distal demonstrative (+ Modifier)
Proximal demonstrative (+ Modifier)
Stressed pronouns + Gesture
Stressed pronouns
Unstressed pronouns
Zeros

This list can be compared with the discussion under the heading "Sufficiently Identifying Language" in chapter 8 above. The term *accessibility* is, of course, used in a way that is very different from its use in the present work. Although Ariel mentions different degrees of "memory availability" (e.g., 1991, p. 444), what is really involved here is the nature of the language necessary to make a shared referent identifiable in a given context. In that light, her discussion is a valuable extension of chapter 8 but could profit from an application to conversational examples.

The other recent way of viewing partially similar material is the "givenness" hierarchy described by Jeanette Gundel, Nancy Hedberg, and Ron Zacharski (1993). They list the following types of expressions, ranging from the least to the most "given." They include the generalized examples on the right which help to clarify labels that are less than optimally mnemonic:

Type identifiable	a N
Referential	this N
Uniquely identifiable	the N
Familiar	that N
Activated	that, this, this N
In focus	it

From the present perspective it appears that what is presented as a single dimension is actually a conflation of activation cost, identifiability, and the functioning of demonstratives. There is a recognition that identifiability is not the same as activation cost, but that problem is solved, not by separating the two dimensions, but instead by appealing to Paul Grice's maxim of quantity (Grice 1975). The validity of that appeal is something I would question, but in any case it would appear that Gundel, Hedberg, and Zacharski, along with Ariel, have forced into a single dimension several aspects of discourse that it would be more profitable to keep apart.

Grammar as Mental-Processing Instructions

Of all the work done in this area, that of Talmy Givón comes closest in spirit to what has been set forth in the present work. It is gratifying to

find that his work has led to understandings that in many ways coincide with or complement those discussed here. We have been aiming at the same target, and if some of his shots seem from the present perspective to have been near misses, they illustrate well the diversity of interpretations that language allows. His recent thinking in the area of information flow was set forth in Givón (1990), where chapter 20 presents an especially useful summary for comparison.

Givón has forthrightly treated language and the mind as inseparably linked, each giving fundamental insights into the other. His perspective on their relationship views grammar (specifically, morphemes and syntactic constructions) as a set of "mental processing instructions . . . designed to trigger specific mental operations in the mind of the speech receiver. . . . These mental operations," he suggests, "involve two well known cognitive domains: (a) attentional activation (b) search in memory storage." It is not obvious why it is necessary to separate activation from memory search, and in fact Givón mentions that "the two may seem coupled or even non-distinct" (Givón 1990, pp. 893–94). One thing he has in mind is the obvious fact that, for the listener, the activation of an idea that is new and unshared cannot involve the reactivation of an idea that is already present in the listener's memory. In such a case the listener's activation takes place without a memory search, although the listener still needs guidance from the speaker in placing the new idea with relation to other, already shared knowledge.

We have here one consequence of the fact that Givón's discussion "is formulated in terms of the speech receiver's . , . perspective. This perspective is adopted for reasons of presentation, and does not prejudge the exact nature of the (at least in part isomorphic) mental processes that take place in the mind of the speech initiator" (p. 895). My own prejudice has been to describe information flow from the perspective of the language producer, who is by definition the person responsible for the form the language takes. I have tried to emphasize, however, how important it is to realize that the speaker's mind necessarily includes a dynamic model of what is happening in the mind of the listener.

Givón makes considerable use of the file metaphor, though he properly notes that such metaphors "tend to be more concrete than their intended mental referents" (p. 895). He says, for example, that "the grammar of referential coherence . . . is about identifying and activating the locations ('files,' 'nodes') where verbally-coded text is stored in episodic memory. The nominal referents-topics seme as 'file labels,' they are used to access ('activate') the storage locations where incoming information is to be 'filed' " (p. 894). I take this to mean that when, for example, an idea that might be verbalized as *Larry* is activated through the use of that word, whatever might be said about Larry will then be assimilated in its proper mental location. It is worth noting that Givón uses the term *referent* for a piece of language (for example, the word *Larry),* and not (as in this book) for the *idea* that may be activated by such a word. What I am calling a referent (the *idea* of a person or object) is apparently equivalent to GivÓn's *storage locution,* file, or *node*.

When Givón says that "verbally-coded text is stored in episodic memory," as in the quotation above, he is aware that it is not language itself that is stored. Elsewhere he points out that "grammatical clues in discourse processing decay rapidly after the message has been decoded, . . . and are thus *not* stored in episodic memory" (p. 940). Although he says that "something like a mental proposition, under whatever guise, is the basic unit of mental information storage" (p. *896),* what he means is that "something analogous to the clause, minus its grammatical form, must be the basic unit of information processing in the mind" (pers. com.). In my terms this basic unit is what I have been calling an idea, most commonly an idea of an event or state, which, when it is verbalized, is likely to take the form of a clause. Given this recognition that information is not stored in verbal form, there remains a problem in understanding the nature of "text-based searches in episodic memory" (p. 941). Although Givón is apparently not suggesting that language comprehenders are literally searching through stored text as such, one wishes that the distinction between verbal and nonverbal storage were more clearly spelled out.

For Givón the notion of *grounding* has considerable importance. Grounding is based in part on the distinction between old and new information, which Givón characterizes as follows: "By 'old' one means 'assumed by the speaker to be accessible to the hearer,' and by 'new' 'assumed by the speaker to be inaccessible to the hearer"' (p. 897). Elsewhere he speaks of old information as predictable, redundant, or topical. Noting that "propositions (or clauses) in coherent discourse . . . tend to be informational hybrids, carrying both old and new information," (p. 898) he goes on to suggest that "the chunks of old, redundant ('topical') information in the clause serve to *ground* the new information to the already-stored old information. Cognitively, they furnish the *address* or *label* for the *storage* locus ('file') in the episodic memory" (p. 899). I hope to have shown, of course,. that the given-new distinction needs to be characterized in terms of consciousness. It should also be noted that Givón's view of grounding differs substantially from that made familiar by Paul Hopper (1979). Nevertheless, there is no arguing with the assertion that speakers include old (or given) information in their clauses as a background for whatever is presented as new.

More problematic, in my view, is Givón's notion of what he calls *topicality.* His development of this notion arose from an understandable dissatisfaction with the variety of ways in which the term *topic* had been used by different investigators (Givón 1983, p. 5). In an attempt to deal with the notion of a topic more effectively, and specifically in order to study

the effect of topicality on the way a referent may be verbalized, he developed several ways of measuring the topicality of a particular referent in a particular context, or at least of finding measures he hoped would correlate with experimental findings on mental processing. These measures were ultimately viewed as ways of specifying two distinguishable components of topicality, which he calls *referential accessibility* and *thematic importance* (Givón 1990, pp. 907–8).

Referential accessibility (or continuity) was seen as measurable in terms of (a) *referential distance* (the number of clauses from the last occurrence of the same referent in the preceding discourse); (b) *switch reference* (whether the preceding clause does or does not have the same referent as an argument); and *(c) potential interference* (the number of semantically compatible referents within the preceding one or two clauses). Of these measures, referential distance became the one most often used in particular studies by Givón and others. Thematic importance was measured in terms of (a) *topic persistence* (the number of times the referent persists as argument in the subsequent ten (earlier three) clauses following the current clause); and (b) *overall frequency* (the total number of times the same referent appears as clausal argument in the discourse). Here it was topic persistence that was most often employed.

The reason for wanting to establish the topicality of a referent in a context was to discover the influence its degree of topicality might have on the way it was verbalized. According to the pattern discovered for English (Givón 1990, p. 913), referents verbalized with unstressed pronouns were found to have a mean referential distance measure of 1. That is, the same referent usually appeared in the immediately preceding clause. Keferents verbalized with stressed pronouns were found to have a mean referential distance of 2.5. That is, the same referent appeared, on the average, two and a half clauses earlier. Referents verbalized with definite nouns were found to have a mean referential distance of 7, but deviation from this mean was so great that the figure could be regarded as meaningless. (For example, while 25 percent of the instances had a referential distance of 1, 40 percent of them had a referential distance of 20 or more.) Finally, so-called left-dislocated definite nouns were found to have a more reliable mean referential distance of 15. Left-dislocation included examples like the following, which must have consisted of two intonation units (Givón 1983, p. 349):

(32) my dad, all he ever did was farm and ranch

Thus, there appeared to be a kind of hierarchy which Givón (1990, p. 913) characterized in terms of, (a) "continuing topics" being coded with "minimal-gap devices" such as unstressed pronouns; (b) "non-continuing topics with anaphoric antecedence within 2–3 clauses back" being coded with "small-gap devices" such as stressed pronouns; and (c) "non-continuing topics with relatively distant anaphoric antecedence" being coded with "long-gap devices" such as left-dislocated definite nouns. These various linguistic devices were seen as instructions to the listener concerning the manner in which the referent in question should be processed.

How would these measures be interpreted from the perspective of the present work? We can consider first the matter of referential distance — the number of clauses between a certain referent and an earlier occurrence of the same referent. This measure can be seen as a rough reflection of activation cost. A referent that was already present in the immediately preceding clause (better, intonation unit) would usually be *given* in the current one. Hence, its expression with an unstressed pronoun would be expected. Perhaps one could regard such a pronoun as an instruction to the listener to interpret the referent as given. It would be misleading, I believe, to interpret the pronoun as a signal of maximum topicality, since the latter term would not be an appropriate way of labeling givenness. What, then, of referents that have a mean referential distance of 2.5 and are verbalized with stressed pronouns? Stressed pronouns, we have seen, usually express contrastiveness, though occasionally they express accessible, noncontrastive referents. Each instance would have to be examined for such properties before one could arrive at any firm conclusions, but it would not be surprising to find that many contrastive referents, if that is what most of them were, were separated from their antecedents by two or three clauses. Keferents that exhibit a large referential distance present a mixed bag from the point of view of activation cost, most of them being either accessible or new. The difference would depend on whether the referent was mentioned at all in the preceding discourse or whether it was being introduced into the discourse for the first time. Since Givón deliberately limited his "look-back'' to twenty clauses, his data would not distinguish accessibility from newness in many cases. The fact that definite nouns show no consistent trend with relation to referential distance, being scattered fairly evenly across the range from 1 to 20 plus, reflects the fact that definiteness — or better, identifiability — is independent of activation cost.

In brief, the "topicality" that is measured roughly by referential distance is largely equatable with activation cost. Unstressed pronouns are usually unambiguous expressions of givenness, stressed pronouns may express either contrastive given referents or accessible referents, and the devotion of an entire intonation unit to an isolated referent (as in "left-'dislocation") may be associated with either accessibility or newness. Identifiability, expressed by definite nouns, is another matter.

To turn to what Givón calls thematic importance, although it may be

measured with either topic persistence or overall frequency it has been the topic persistence measure that has usually been applied—the number of times the referent persists as an argument in the subsequent ten (formerly three) clauses following the current clause. What is being measured here is evidently what I called referential importance in chapter 7, where I concurred with Givón's suggestion that one way of determining such importance operationally is to count the number of occurrences of the referent within the relevant stretch of discourse (cf. Wright and Givón 1987).

It is interesting to see how Givón relates topicality and subjecthood. "The quantified study of the topicality of grammatical subjects and objects in connected discourse" shows that "the subject is consistently more topical than the direct object, and the direct object more topical than the indirect object" (Givón 1990, p. 901). In terms of the present work, Givón's hierarchy would say that subjects are most often given and of primary importance, that direct objects rank somewhat lower on the scales of activation cost and importance, and that all other roles rank lower still. The strong correlation between givenness and subjecthood is beyond doubt, and it is true that most subjects are of either primary or secondary importance. These two properties are effects of the role of subjects as the grammaticized expression of starting points. Givón's model would be more congruent with the present one if topic were equated with starting point, referential accessibility with activation cost, and thematic importance with what i am calling referential importance. It would be necessary, however, to recognize the separate status of all three, being alert to the various ways in which they interact as well as the functional reasons for such interactions.

I believe it is not reading too much into Givón's work to suggest that he has, following a somewhat different path, recognized both the light subject constraint and the one new idea constraint. When, for example, he states as a "general principle" that "only one file is open at any given time" (1990, p. 939), I believe he could be translated as saying that each clause has a single starting point, which, as we have seen, is most likely to be "highly topical" or, in my terms, "light." And when he states that "a clause in connected discourse tends to contain only one chunk of new information" (p. 898), he is obviously talking about the one new idea constraint—as he puts it, "an expression of some *cognitive limit* on the processing rate of new information." Less clear is the question of whether he has created a place for what I have been calling accessible information? as might be the case when he says that "grammar-guided discourse processing seems to involve covert attention" (p. 939). Most important, however, is Givón's recognition that further understanding depends on a broader vision of our task, as when he writes of bringing together, "within a single coherent framework, facts from the hitherto disparate domains of grammar, discourse and cognitive psychology. All other things being equal, this is the type of increasing scope of coherence that one strives for in science" (p. 941).

Summary

From the perspective of this book, all the lines of research discussed in this chapter could profit from more clearly differentiating the roles of activation cost, contrastiveness, starting points, referential importance, identifiability, newsworthiness, and no doubt other discourse functions. Except for Halliday, these traditions have tended to lump together diverse functions under unitary labels such as "communicative dynamism," "givenness," "assumed familiarity," "accessibility," or "topicality." I believe, of course, that it is also essential to recognize the central role of consciousness, no longer characterizing given information, for example, as "known," "retrievable," "predictable," "recoverable," "familiar," "accessible," or the like. It would help to give a place to semiactive consciousness and its relevance to what I have called accessible information. Finally, although several of the researchers discussed in this chapter have explicitly recognized the importance of working with natural discourse, and especially conversations, all of them have in practice relied on mixtures of spoken and written, real and constructed data. Despite these differences, all of us have been groping toward much the same goal, and some convergence seems gradually to be emerging. A deeper and wider-ranging survey of the sort sketched in this chapter will undoubtedly shed useful light, not just on the subject matter itself, but also on the trajectories and discontinuities that have characterized the recent histories of linguistics and related disciplines.