

Chapter 9

Inferring What is Meant

Herbert H. Clark

How do people comprehend what others say? The answer to this question depends on what one conceives 'comprehension' to be. Many investigators (see Fodor, 1975; Tanenhaus, Carroll, and Bever, 1976) seem to view it as the process by which listeners arrive at the literal or direct meaning of a sentence (as described by generative linguists). This I will call the *independence view*, for it is concerned with only those processes that apply to a sentence independent of its context. On the face of it the independence view seems overly narrow. It excludes much of what 'comprehension' is ordinarily thought to include—for example, the identification of referents. In an everyday conversation, if David said *He's crazy* and Nancy did not identify who 'he' was, she would ordinarily be said not to have understood David completely. Other investigators (see Bransford and Johnson, 1973; Bransford and McCarrell, 1974) view comprehension as the process by which listeners, relying on real world knowledge, build elaborate mental edifices for the situation a sentence describes. This I will call the *constructivist view*. Yet in critical ways this view seems too inclusive. On hearing *He's crazy*, Nancy might adventitiously be reminded of her mad Uncle Harry, which in the constructivist view would become part of her understanding of the utterance. But in the ordinary sense of 'comprehension', she wouldn't be said *not* to have understood David if she had not had this tangential thought.

In this chapter I will argue instead for an *intentional view* of comprehension, one that lies between the independence and constructivist views. In it comprehension is conceived to be the process by which people arrive at the interpretation the speaker *intended* them to grasp for that utterance in that context (see Grice, 1957; Schiffer, 1972; Bennett, 1976). Unlike the independence view, this view requires listeners to draw inferences that go well beyond the literal or direct meaning of a sentence. It insists that Nancy should identify who David is referring to in *He's crazy*. But unlike the constructivist view, it limits the inferences to those that listeners judge the speaker intended them to draw. It excludes stray thoughts about mad uncles. In this view the speaker's intentions

are critical, but they can ever only be inferred. My aim in this chapter, then, is (a) to illustrate when and how people infer what is meant and (b) from there suggest some general ways for formulating a more satisfactory theory of comprehension.

INTENDED INTERPRETATIONS

In the intentional view of comprehension there is a distinction between the literal, or direct, meaning of a sentence and the intended interpretation of that sentence as uttered in a particular context. But what is the distinction? This question is not easy—indeed answering it is a major goal of philosophers interested in this view. Yet the essential distinction should become clear with some examples.

Take *He's crazy*. Considered in isolation, it presupposes that there is a male entity referred to by the definite noun phrase *he*, and it asserts that this entity is insane. This is roughly its literal or direct meaning. But when Nancy hears David say *He's crazy* in an ordinary conversation, she does not want to know merely that there *exists* a male entity to which *he* refers. She wants to know which of many conceivable entities it actually *does* refer to. For Nancy this poses a problem. There is nothing in the sentence *He's crazy* to tie *he* to the actual person it refers to. Nancy must go beyond the literal meaning and infer the connection, and that will not always be easy.

Or take *In Italy Margaret fell down and broke her arm*. In this sentence, *and* conjoins two clauses—*Margaret fell down* and *Margaret broke her arm*—with the result that both events are asserted to be true. Yet in most circumstances people would also take it that Margaret's fall was being said to be the *cause* of the break. People not drawing this inference would normally be thought not to have understood the utterance completely. Note that the causal interpretation is not a necessary one. The sentence could have meant merely that Margaret fell down at one time and broke her arm at another, in analogy with *In Italy Margaret visited Milan and met people in Rome*. So for this utterance people have to go beyond the literal meaning and infer the causal interpretation that they think is intended.

Or take *San Francisco always votes Democratic*. Its literal interpretation does not make sense. A city is an inanimate physical location that cannot go to its local polling place and cast votes. One way around this anomaly is to assume that *San Francisco*, in its literal meaning, is elliptical for *one or more persons connected with San Francisco*. But in ordinary conversations even this meaning will not do. In the right context *San Francisco* will be taken to mean 'a majority of people who actually vote in the election of interest', which is far more specific than the literal meaning. In other circumstances it will be taken to mean 'the mayor of San Francisco', as in *On the Council of Mayors San Francisco always votes Democratic*, 'a majority of the people on the city council', as in *Of the three*

city councils San Francisco always votes Democratic, and so on. People are not satisfied with the literal meaning of *San Francisco*. In each context they try to find the specific category it is meant to pick out, and that requires inferences.

Or take this short conversation:

Alex: Are Levelt and d'Arcais psychologists?

Barbara: Is the Pope Catholic?

In answer to Alex, Barbara asked a question, but she did not mean it to be taken as merely a question. She was using it to say 'Emphatically yes' to Alex's question. He was meant to go beyond the literal meaning ('I request you to tell me whether or not the Pope is Catholic') to infer the interpretation actually intended. If he had failed to do this, he would ordinarily be said to have misunderstood what Barbara said.

In the everyday sense of 'comprehension' or 'understanding', then, people try to get at what they were meant to understand, and this ordinarily forces them to draw inferences that go beyond literal meaning. But what are these inferences? And how are they drawn? These questions are best approached, I will argue, with comprehension treated as a form of problem-solving. In the next section I will outline such a treatment, and in the following three sections, I will take up its consequences in three domains I have been especially interested in—indirect requests, definite noun phrases, and shorthand expressions.

COMPREHENSION AS PROBLEM-SOLVING

In getting at the intended interpretation of an utterance, listeners have a problem to solve: What did the speaker mean by what he said? The sentence uttered provides only some clues to the solution. Other clues must be sought in the physical and linguistic context of the utterance, the speaker's assumed beliefs, and all kinds of general knowledge. These other clues change from moment to moment, with some relevant one instant and others relevant the next. There is no algorithmic method by which listeners could anticipate all possible combinations of clues. They can only collect and weigh the evidence available and, by manoeuvres common to other types of problem-solving, infer what the speaker must have meant.

One view of this process is outlined in Table 9.1. First, there is the goal: find the intended interpretation. Next, there is the information, or 'data base', that is relevant to achieving the goal. This includes the sentence uttered, the time, place, and other circumstances of the utterance, the speaker's assumed beliefs, and general background knowledge about objects, states, and events that do or can exist in the real world. Third, there are the special assumptions that people make about the act of communication. These are set apart here as boundary conditions. As I will note later, these assumptions take the form of tacit

Table 9.1 Solving for the intended interpretation

Goal: What is the intended interpretation?

Data base: The sentence uttered; the time, place and circumstances of the utterance; the speaker's beliefs about the listener; general knowledge.

Boundary conditions: Various tacit agreements between speaker and listener about how language is to be used.

Mental operations:

1. Build a candidate interpretation.
 2. Test the candidate interpretation against the boundary conditions.
 3. If it passes all the tests, accept it as the intended meaning.
- Otherwise begin at 1 again.
-

agreements speakers and listeners have with each other about how language is used. Finally, there are the three general mental operations people call on in solving problems. The first builds candidate interpretations; the second weighs the evidence for and against them; and the third registers one of them, if it passes muster, as the intended interpretation, the solution to the problem. I should stress that I have no special commitment to this particular way of characterizing the problem-solving, but it allows me to make a start on how the process might work.

Treating comprehension as problem-solving does not mean that the process is conscious. Most strategies listeners have available for building, testing, and registering interpretations are quick and efficient and carried out without awareness. Much of the process can probably never be introspected about. This is no different from problem-solving in arithmetic, algebra, geometry, and other domains, where most mental operations, especially on simple problems, are done without awareness.

Tacit agreements

Central to this characterization of problem-solving are the tacit agreements, the boundary conditions through which listeners weigh, accept, and reject possible interpretations. The notion of tacit agreement comes from Grice (1975), who proposed that in order to communicate successfully, people tacitly agree to cooperate with each other, to adhere to what he called the cooperative principle. In particular, it is agreed that speakers will ordinarily try to follow four 'maxims': (a) be informative, (b) be truthful, (c) be relevant, and (d) be clear. Thus, in composing utterances speakers will try to conform to these maxims, and in interpreting those utterances listeners will assume that speakers are trying to conform to them.

The cooperative principle and its four maxims, Grice has demonstrated, are critical in deciding what speakers mean. Consider the maxim 'Be truthful'. If

people did not believe the things they asserted were true—if they did not have sufficient evidence for them—communication would be chaotic. Listeners would have no way of distinguishing what speakers believed from what they did not believe, and factual information could not be communicated. The maxim seems fundamental to communication. But it can also be used for conveying information indirectly. Imagine that after seeing an appalling production of Hamlet, Cleo says to Tony *Wasn't that an exquisite production!* He realizes she knows that he thinks it was a bad production and that he knows she thinks it was a bad production. Yet she appears to have said just the opposite, a blatant untruth. Even though she is flouting the maxim 'Be truthful', however, Tony has no reason to think she is not trying to be cooperative, so she must not mean what she said to be taken literally. She must be *implicating* that she is being sarcastic, that she means 'What a terrible production!' Thus, by reference to these tacit agreements, Tony has gone beyond the literal meaning of what Cleo said to infer what she meant. *Implicatures* like this, Grice has shown, can arise on the basis of the other three maxims as well.

In discussing the cooperative principle, however, Grice was rather sketchy. He did not elaborate on how it would apply to anything but a few well-chosen examples. So far it forms only the skeleton of a theory of communication, and it needs to be filled out before it can do the work it was meant to do. For each area I will take up, therefore, I will try to be more specific about what tacit agreements might be operating and how they would apply.

In summary, comprehension will be treated as a form of problem-solving. Listeners are faced with the problem of discovering what the speaker intended in uttering what he did. To achieve a solution, they sift and weigh all the evidence at their disposal—sentence, context, general knowledge, and judgments about the speaker's beliefs. They inspect the evidence particularly closely in the light of certain tacit agreements they share with the speaker about how language is to be used. These agreements lead them to apply specialized strategies that are particularly quick and efficient—that lead to the solution with the least amount of time and effort. Nevertheless, solving for the intended interpretation is not always easy. It takes measurable time and effort and leads to other behavioural consequences. In the next three sections I will consider some of these consequences for three types of expressions—indirect requests, definite noun phrases, and shorthand expressions. I have chosen these because they range in size from complete utterances (indirect requests) to single words (shorthand expressions) and I want to show how inferential strategies are required no matter what size of expression is considered.

INDIRECT REQUESTS

Requests are a type of interpersonal negotiation in which one person tries to get another to do something, like pass the salt. English has a construction especially

designed for this purpose, the imperative. It is used for making direct requests, as in *Please pass the salt*. But requests can be made in other ways too, depending on the context, as in *Can you pass the salt?*, *Can you reach the salt?*, *I want the salt*, *If you don't pass the salt this minute, I'll scream*, and *This soup needs salt*. Since all of these use constructions other than the imperative, they are often called indirect requests. But how do people come to regard such utterances as requests? How do they arrive at their intended interpretations?

According to Austin (1962) and Searle (1969, 1975), people have tacit agreements about what constitutes a proper, genuine, or felicitous request. To be proper it must satisfy four 'felicity' conditions:

1. *Preparatory condition*. The speaker believes the listener *is able to* carry out the requested act.
2. *Sincerity condition*. The speaker *wants* the listener to carry out the requested act.
3. *Propositional content*. The speaker predicates a *future act* (the one being requested) of the listener.
4. *Essential condition*. The speaker counts his utterance as an attempt to get the listener to carry out the requested act.

People tacitly agree to make requests only when these four conditions are fulfilled. So when any one condition is not satisfied, requests go wrong. For example, Julia could not felicitously ask Ned to fly to the moon, because that is not something she believes he is able to do. The preparatory condition would not be satisfied. Requests go wrong in other ways when other conditions are not satisfied. Thus, like all other interpersonal negotiations, requests are made against a background of tacit agreements.

All this holds for indirect requests too, but they are identified as requests via the felicity conditions themselves. According to Searle (1975; see also Gordon and Lakoff, 1971) indirect requests fall into four major classes. First, speakers can ask whether, or state that, the preparatory condition holds—that the listener is able to do the requested act. In the right context one can say *Is it possible for you to pass the salt?* or *You could pass the salt* and it will be taken as a request. Second, speakers can state that the sincerity condition is satisfied—that they want the requested act carried out. This leads to such indirect requests as *I want you to pass the salt* and *I would appreciate your passing the salt*. Third, speakers can ask whether, or state that, the propositional content condition is satisfied—that the listener will do the requested act in the future. This class includes such indirect requests as *Won't you pass the salt?* and *You will pass the salt right this minute*. And finally, speakers can ask whether, or state that, there are compelling reasons for the requested act being done, as in *Why don't you pass the salt?* and *This soup needs salt*. These really belong to the second class,

those having to do with the sincerity condition, since speakers should not request something unless they have good reasons for wanting it done and think there are no compelling reasons for it not being done. In short, speakers can *imply* they are requesting something merely by suggesting that one of the felicity conditions for that request is fulfilled. This is a powerful and useful method.

Central to this method is the notion that speakers can use one interpretation to imply another. Imagine that Helen tells Margaret *This soup needs salt*. Helen means her utterance to be taken literally—she is asserting that the soup really does need salt—but she is using this interpretation to do something further—request Margaret to pass the salt. That is, she means to convey both the assertion and the request, with the assertion being only a means to the request. So it is not that the construction *This soup needs salt* is being used as a request, but rather that the assertion that the soup needs salt—the interpretation of that construction—is being used to make a request. Listeners are meant to register both interpretations.

Chains of interpretations like this can become very long indeed. Imagine that George says to his son Ken, *Haven't you forgotten to clean up your room?* In the right context, this utterance will form a chain of interpretations something like this:

1. *Yes/no question*. Is it or is it not the case that you have forgotten to clean up your room?
2. *Assertion*. You have forgotten to clean up your room.
3. *Reprimand*. You haven't cleaned up your room, as you were supposed to.
4. *Assertion*. I want you to clean up your room.
5. *Request*. Clean up your room.

The direct interpretation, a question, is used to make an assertion, the assertion to make a reprimand, the reprimand to make a second assertion, and the second assertion, via the sincerity condition on felicitous requests, to make a request. Ken is meant to register all of these interpretations at once, but in relation to each other as the chain specifies. For example, if he says *No, I've already done it*, he has answered the first question in such a way as to nullify the force of the remaining four interpretations. Or if he says *But I wasn't supposed to*, he denies the reprimand, nullifying the prior assertion and question, and questioning the basis for the following assertion and request. In short, an utterance like this has an initial interpretation, zero or more intermediate interpretations, and a final interpretation. While the ultimate reason it is uttered may be to convey the final interpretation, it uses the prior interpretations in a particular sequence to do this.

Empirical Consequences

How do listeners arrive at the intended interpretation of indirect requests? The facts just considered suggest an approximate first answer, which I will present as a process model:

- Step 1.* Compute the direct interpretation of the utterance.
- Step 2.* Decide if the interpretations computed so far are what were intended. Are there sufficient and plausible reasons for the speaker to have intended these interpretations alone in this context?
- Step 3a.* If yes, proceed to Step 4.
- Step 3b.* If no, use the immediately preceding interpretation to compute an additional interpretation by way of the tacit agreements on speech acts. Then return to Step 2.
- Step 4.* Utilize the utterance on the basis of its collection of interpretations and assume that the final interpretation is the ultimate reason for the utterance.

This is one form of the problem-solving process presented earlier, for it builds candidate interpretations, tests them against the boundary conditions, and accepts them when they fit. But in its more specific form it has several testable implications. I will consider only four of them.

The first implication concerns the equivalence-of-interpretations assumption: indirect interpretations should have the same properties as the equivalent direct interpretations. The request interpretation of *Won't you go home?*, for example, should exhibit the same properties as that of *Please go home*, despite the fact that *Won't you go home?* conveys the interpretation only indirectly. This consequence has been tested in an experiment I carried out with Lucy (Clark and Lucy, 1975). In it people were shown a series of requests each accompanied by a picture. For each request they were to decide as quickly as possible whether or not the act being requested had been carried out in what was depicted in the picture. Several of the requests we used were direct, like *Please open the door*, while the rest were indirect, like *I would love to see the door opened*. We looked at whether the indirect requests showed the same pattern of verification times as the direct requests.

Requests like these can have either a positive or a negative force. The direct requests *Please open the door* and *Please don't open the door*, for example, are related as positive to negative, and so are the indirect requests *Can you open the door?* and *Must you open the door?*, *Why not open the door?* and *Why open the door?*, *I would love to see the door opened* and *I would hate to see the door opened*, and other such pairs. This contrast allowed us to make several crucial comparisons, for much was already known about verification times for positive and negative assertions. In a study by Clark and Chase (1972), for example, the

time it took people to judge affirmative and negative assertions true or false of accompanying pictures were as follows:

True Affirmative. Star is above plus.	1810 msec
False Affirmative. Plus is above star.	1997 msec
True Negative. Plus is not above star.	2682 msec
False Negative. Star is not above plus.	2495 msec

These findings are typical (see Clark, 1976). Affirmative sentences are judged faster than negative sentences—here by 685 msec. And for affirmatives, true judgments are made faster than false ones, while for negatives, false ones are made faster than true ones. We thought that if the request interpretations of, for example, *Can you open the door?* and *Must you open the door?* are positive and negative in the same way, they should yield a similar pattern of latencies.

Indeed, they did. The following latencies are for only two of the ten pairs of requests we used, but they exhibit the main findings:

Direct requests

<i>True.</i> Please open the door.	1213 msec
<i>False.</i> Please open the door.	1610 msec
<i>True.</i> Please do not open the door.	1799 msec
<i>False.</i> Please do not open the door.	1644 msec

Indirect requests

<i>True.</i> Can you open the door?	1473 msec
<i>False.</i> Can you open the door?	1990 msec
<i>True.</i> Must you open the door?	2082 msec
<i>False.</i> Must you open the door?	1810 msec

First, the indirect requests exhibited virtually the same pattern of latencies as the direct requests. Second, both exhibited the main features common to studies of affirmative and negative assertions. Affirmatives were faster than negatives, and for affirmatives, true was faster than false, and for negatives, false was faster than true. Thus, these findings constitute evidence for one major consequence of the model outlined earlier: a request interpretation has certain properties no matter whether it is conveyed directly or indirectly.

A second implication of the model comes from what might be called the order-of-interpretation assumption: most indirect interpretations are computed on the basis of logically prior direct interpretations. For Helen's *This soup needs salt*, for example, Margaret computes the direct interpretation—Helen is asserting that the soup needs salt—and uses it to compute the indirect interpretation—that Helen is requesting Margaret to pass the salt. What this

implies is that the longer it takes to compute the initial interpretation, the longer it should take to arrive at the final interpretation.

This implication was also tested in our study, at least for certain types of sentences. Two of our pairs of indirect requests were of this type:

- (1) I'll be very happy if you open the door. (Positive)
I'll be very sad if you open the door. (Negative)
- (2) I'll be very sad unless you open the door. (Positive)
I'll be very happy unless you open the door. (Negative)

In each pair the first request is positive and the second negative, but the way in which people should arrive at these interpretations for the two pairs is different. In (1) they should first compute a direct interpretation, making use of *if*, and use that to compute the indirect interpretation. In (2), the process should be the same except that the computation of the direct interpretation depends on *unless*. But *unless* is inherently negative—it means something like 'only if not'—and ought to take longer to make use of than *if*. Thus, if the model is right and people compute indirect interpretations on the basis of direct interpretations, they should take longer on (2) than on (1). They did. Pair (2) took about one second longer in our task than did Pair (1).

One caution here. In these two pairs it seems obvious that listeners must compute the direct before the indirect interpretations. How could they know whether the request was positive or negative unless they determined how the speaker's happiness was contingent on the door's position? In other instances, the point is not so clear. Some indirect requests, like *Can you open the door?* and *Why not open the door?*, may be idioms analogous to the 'die' interpretation of *kick the bucket* (Sadock, 1972, 1974). That is, the request interpretation of *Can you open the door?* is one of its two direct interpretations and is *not* dependent on its other direct interpretation—'I ask you whether or not you are able to open the door'. If so, the request interpretations of *Can you open the door?* and *Please open the door* would be dealt with alike in the model presented earlier, for both would be treated as direct interpretations. But the issue is even more complex. Indirect requests, like other constructions, probably lie on a continuum of idiomaticity (see Bolinger, 1975, pp. 99–107). Some may be dealt with directly, others indirectly, and still others directly or indirectly depending on the circumstances. At present very little is known about idiomaticity in comprehension, although the issue seems to be crucial.

A third implication of the model follows from the assumption that listeners register the whole chain of interpretations for an utterance—from its direct to its ultimate interpretation. If they register all these interpretations during comprehension, later, when asked to remember what had been said, they should be able to recognize not just the direct or the ultimate indirect interpretation, but the whole chain. Consider these three sentences, whose ultimate

interpretations in context are on the right:

- (3) The food is on the table. (Invitation)
- (4) The food is on the table (Assertion)
- (5) Please take the food on the table. (Invitation)

In their direct interpretations (3) is the same as (4), while in their ultimate indirect interpretations (3) is the same as (5). If listeners originally heard (3), they should therefore be able to distinguish it from both (4) and (5), since neither one contains the whole chain of interpretations of (3).

This point was tested in an experiment by Jarvella and Collas (1974). In it people were asked to take the part of an actor reading a script aloud. Later, in reading a second script aloud they were stopped at various points and asked whether or not the sentence just read was the same as one in the first script. One sentence in the first script, for example, was (3), and on the second script different people were asked about (3), (4), or (5). Averaged over many analogous examples, the findings were clear. People readily recognized (3) in the second script as being the same as (3) on the first. But they were reliably less willing to say that (4) or (5) on the second script was the same as (3). The percentages of 'same' judgments for (3), (4), and (5) on the second script were 89%, 77%, and 69%, respectively. Yet people recognized the content of (4) and (5) much better than chance. Faced with the invitation *The food is on the chair* in the second script, people were willing to say 'same' only 7% of the time. In short, people appear to retain the chain of interpretations behind indirect requests like (3), implying they compute and register the whole chain.

The fourth implication follows from the same assumption as the third—listeners register the chain of interpretations, not just the ultimate one—and it too concerns memory. Consider (6) and (7) (from Keenan, MacWhinney, and Mayhew, 1977):

- (6) I think there are two fundamental tasks in this study.
- (7) I think you've made a fundamental error in this study.

In their direct interpretations (6) and (7) are both assertions, but in the right context (7) would be interpreted indirectly as a reprimand. That is, (6) has a single interpretation in its chain, while (7) has at least two. If people comprehend and register the whole chain, they will have essentially one memory code for (6) but two for (7). With the two codes for (7), they have a better chance of recognizing it. They even have a better chance of distinguishing it from close paraphrases of its direct meaning, for many of these paraphrases will not lead to the right indirect interpretations.

These two predictions were tested indirectly in a study by Keenan and her colleagues. In it people who had attended a midday discussion were

unexpectedly asked one or two days later to judge whether or not certain sentences had been uttered during that discussion. This recognition test included sentences that had actually occurred, paraphrases of sentences that had actually occurred, and plausible sentences that had not occurred. Of the sentences in the first two categories, half had what Keenan *et al.* called 'high interactional content'. Like (7) they contained 'information about the speaker's intentions, beliefs, and his relations with the listener'. The other half, like (6), had 'low interactional content'. From the examples Keenan *et al.* presented, it appears that the high interactional content sentences always had both a direct and an indirect interpretation, while the low interactional content sentences normally had only a direct interpretation. So this study provides a rough test of the fourth implication.

Keenan *et al.*'s findings were clear. The sentences with high interactional content—those with a chain of interpretations—were recognized far more readily than those with low interactional content—those with only a single interpretation. Indeed, the high interactional sentences were also more accurately distinguished from their paraphrases. As a control Keenan *et al.* asked people not at the lunchtime discussion to study the same high and low interactional sentences as lists of sentences and then take the same recognition test. Because the sentences were not heard in context, they were all probably interpreted more or less literally, hence were all essentially 'low interactional content' sentences. As expected, the two types of sentences were recognized about equally often. So it appears that the high interactional sentences were recognized better in context because they took on additional interpretations that made them easier to tie into the original conversation and easier to distinguish from literal paraphrases.

Indirect requests—and indirect interpretations in general—suggest that intentions are central to the comprehension process. On encountering an utterance, people ordinarily ask, 'Why did the speaker say that here and now?' and are not satisfied until they find a plausible reason. This often means computing the direct interpretation, asking whether or not it is sufficient in this context, and computing further indirect interpretations when it is not. The evidence cited here seems not only to fit these ideas, but to require them. Nevertheless, the details of this process are still far from clear.

DEFINITE REFERENCE

When listeners are confronted with a referring expression like *the woman* or *the knife*, they assume not only that such a woman or knife exists in some possible world, but also that the speaker expected them to identify the one he had in mind. Consider the following sequences:

- (8) I met a man and a woman yesterday. The woman was a doctor.
- (9) I went to a wedding yesterday. The woman was a doctor.
- (10) A man fell to the floor murdered. The knife was dropped nearby.

In all three, listeners have to decide what is being referred to by the definite noun phrases. In (8), *the woman* almost certainly refers to the woman mentioned in the first sentence, but even that requires an inference. In the right context it could refer to another woman identifiable by another means. In (9), where there is no woman explicitly mentioned in the first sentence, listeners have a more difficult problem to solve. They must see that since the first sentence mentions a wedding, and since every wedding has a bride, *the woman* probably refers to the bride. In (10), the inference is also rather complicated. Listeners must assume that *the knife* refers to something about the event just mentioned, the man's murder. Since murders usually require weapons, and since knives are possible weapons, *the knife* probably refers to the weapon used in the murder. Definite reference, then, is an example *par excellence* of having to infer what is meant.

Like other expressions, however, definite noun phrases are used within the confines of a tacit agreement. This one, which happens to be part of an agreement Haviland and I have called the *given-new contract* (Clark, 1977; Clark and Haviland, 1974, 1977; Haviland and Clark, 1974), goes something like this:

The speaker agrees to use a definite noun phrase only when he has a specific referent in mind and is confident that the listener can identify it uniquely from its description in the noun phrase.

For example, if Nancy were to say (10) to Jeffrey, she would have to have a specific knife in mind and fully expect Jeffrey to be able to identify it from what she believed he already knew. From Jeffrey's point of view, it is pertinent that Nancy had described the object as a knife, that she had just mentioned a murder, that knives are common murder weapons, and even that she assumed he could identify the knife uniquely from the little information she provided. Indeed, the last piece of knowledge—derived from the tacit agreement itself—is far more critical than one might suppose.

How do listeners solve for the intended referent? One hypothesis is that the problem-solving process in Table 9.1 takes on a specific form something like this:

- Step 1.* Compute the description of the intended referent.
- Step 2.* Search memory for an entity that fits this description and satisfies the criterion that the speaker could expect you to select it uniquely on the basis of this description. If successful, go to Step 4.
- Step 3.* Add the simplest assumption to memory that posits the existence of an entity that fits this description and satisfies the criterion that the speaker could expect you to select it uniquely on the basis of this description. If successful, go to Step 4.
- Step 4.* Identify this entity as the intended referent.

To see how this works, consider *the woman* in (11) and (12):

(11) I met a man and a woman yesterday. The woman was a doctor.

(12) I met two people yesterday. The woman was a doctor.

At Step 1, listeners would compute the description of the intended referent, namely that it is a woman. At Step 2, they would search their memories for such an entity. In (11), they would find the woman just mentioned. That entity satisfies the criterion that it could be selected uniquely on the basis of this description, so listeners could proceed to Step 4 and identify her as the intended referent of *the woman*. In (12), however, no woman has been mentioned, so listeners would have to proceed to Step 3. There they would add a bridging assumption, the simplest one possible, that posits the existence of a woman and still satisfies the criterion of unique selectivity. Without other information they would most likely add this assumption:

(12)a. One of the two people mentioned is a woman and the other is not.

The woman posited here fits the wanted description and could then be selected uniquely. But note that to get uniqueness, (12a) had to specify not just that one of the two people was a woman, but that the other one was not. With (12a) added, listeners would proceed to Step 4 and identify that woman as the intended referent of *the woman*.

The process just described is anything but complete. It doesn't say how listeners search their memories at Step 2. For example, do they try entities just mentioned first? Do they search the memory in a systematic order? Nor does it say when listeners give up searching (Step 2) and begin building bridging assumptions (Step 3). Steps 2 and 3 may work in parallel, with priority given to entities identified in Step 2. Nor does it say what the 'simplest' bridging assumption is, or how listeners go about building such an assumption. The model is incomplete in other ways too. If it has value, then, it is in providing a first examination on which one can build.

Empirical Implications

Yet the proposed model has several empirical implications. First, if listeners ordinarily compute referents in comprehending definite noun phrases, they should realize when they have failed to find what they believe to be the intended referent. Second, if they need to add bridging assumptions in order to compute an intended referent, they should show evidence that they have added them. And third, all of this process takes time. In earlier papers, Haviland and I (Clark and Haviland, 1974, 1977) discussed these and other implications in some detail. Here I will review only some of the pertinent evidence.

The first implication is almost too obvious. Listeners should object to speakers who violate their tacit agreement about definite reference. This agreement can be broken in several ways. For example, imagine that George, on meeting Margaret one day, says *He broke a ski*. Margaret's first reaction would be to say, 'I don't understand. Who is *he*'?. George has violated the part of the agreement covering computability—Margaret does not have enough information to find a plausible referent for *he*. Or imagine George telling her *Ken and Bob went skiing yesterday and he broke a ski*. This time Margaret would object, 'But which one broke the ski—Ken or Bob?'. George has violated that part of the agreement covering uniqueness. Although Margaret can compute a plausible referent for *he*, she cannot do so uniquely. At the same time listeners should realize when they should not be able to compute the intended referent. For example, when Jane overhears George tell Margaret *He broke a ski*, she should not be surprised if she cannot guess who 'he' is. George designed his definite reference for Margaret, and not necessarily for anyone else. With her current information Jane may or may not be able to identify the intended referent.

The second implication is that when bridging assumptions are necessary, listeners should show evidence of having added them. The major evidence for this is introspective. For definite reference in the appropriate circumstances, listeners readily agree that they have made certain bridging assumptions. The point can be made most vividly by considering several examples from a brief taxonomy of definite reference (Clark, 1977). Each example requires a bridging assumption, and merely reading them shows how readily and irresistably we build the right ones. The examples fall into three major types—direct reference, indirect reference to parts, and indirect reference to roles.

With direct reference the bridging assumptions are relatively simple, as illustrated in these four examples:

- (13) I saw a play last night. The play I saw last night was by Tom Stoppard.
- (14) I saw a play last night. It was by Tom Stoppard.
- (15) I saw a play last night. The damn thing was by Thomas Dekker.
- (16) I saw two plays last night. The brilliant one was by Stoppard.

In reading (13) and (14) we do not add any significant bridging assumptions. In (13) the referring expression *the play I saw last night* contains all the information we know about the play mentioned in the first sentence, so it is simple to infer what it refers to. In (14), although *it* contains only a scrap of that information, the inference is almost as easy to make. In (15), however, we are forced to add a significant assumption. Although the epithet *the damn thing* has the same function as *it* in (14)—it refers to the previously mentioned play—it forces us to

add this bridging assumption:

(15)a. The speaker didn't like the play he saw last night.

In (16), *the brilliant one* picks the referent from a class previously mentioned, but we all add an assumption something like this:

(16)a. The play referred to by *the brilliant one* was brilliant, and the other one wasn't.

We seem to form this assumption effortlessly.

For indirect reference to parts, the bridging assumptions are more complicated—more obviously inferred—as in these three examples:

(17) Robert found an old car. The steering wheel had broken off.

(18) Robert found an old car. The radio was still in good shape.

(19) Robert found an old car. The letters F-O-R-D were still clearly visible.

We readily infer in (17) that the car had a steering wheel, in (18) that it had a radio, and in (19) that it had the letters F-O-R-D printed on it. We also infer that these are the objects being referred to by *the steering wheel*, *the radio*, and *the letters F-O-R-D*. But we did not add these bridging assumptions unthinkingly. While every old car has a steering wheel, few have radios, and even fewer have F-O-R-D printed on them. In (18) and (19), then, it was the referring expressions *the radio* and *the letters F-O-R-D* that led us to add the appropriate assumptions. We would never have made such strong assumptions without these goals.

With indirect reference to roles, the bridging assumptions get even more complicated. Consider (20) and (21).

(20) Yesterday there was a killing in Saloon Number Ten. The victim was Wild Bill Hickok.

(21) Yesterday there was a killing on Broadway. The getaway bicycle was later found in Central Park.

In (20) we readily assume that the victim referred to is the person killed in the killing mentioned in the first sentence. Killings necessarily have victims, so this bridging assumption is easy to make. But in (21) we just as readily assume that the killer made a getaway on a bicycle, and that that bicycle is the one being referred to. Yet it is not necessary or even usual for bicycles to play such roles in killings, even on Broadway. Thus, even though the bridging assumption for (21) is worked out from intricate pieces of world knowledge, it is one we make automatically. We are surprised, on reflection, that it is so complex and so indirectly inferred.

Bridging assumptions added in this process show up particularly clearly in experiments on memory. Keenan and Kintsch (see Kintsch, 1974) asked people to read one of these two sequences (and others like them):

- (22) A carelessly discarded burning cigarette started a fire. The fire destroyed many acres of virgin forest.
- (23) A burning cigarette was carelessly discarded. The fire destroyed many acres of virgin forest.

Later both groups were asked to judge whether *The discarded cigarette caused the fire* was true or false of what they had read. In (22), of course, this information is stated directly, so it is no wonder people almost always said 'true'. In (23), however, it is not stated outright. Rather, it is a bridging assumption people need only to identify the referent of *the fire*. But people in this group almost always said 'true' to the test sentence also, as they should have if they had added the assumption while reading the sequence. Yet for (23) one could argue that it was the test sentence, and not the earlier noun phrase *the fire*, that led them to infer that the cigarette had caused the fire. In fact, when people were tested immediately after reading (22) or (23), they were slightly faster at judging the test sentence for (22) than for (23). But if this alternative explanation were correct, they should also have been faster on (22) than on (23) when the test sentence was presented 15 minutes later, and they were not. Apparently, people judged the implicit information in (23) more slowly immediately afterwards, not because the information was not there, but because it was less available for conscious judgments. Thus, as many other memory experiments also show (see Clark and Clark, 1977, Chapter 4), people identify referents, adding the necessary bridging assumptions as they go along.

The third consequence to be considered is that it takes people time to search their memories and decide on the intended referent. This was demonstrated in a study I carried out with C. J. Sengul. We were interested in how long it took people to identify a referent depending on how recently it had been mentioned in the preceding discourse. So we composed a number of sequences, like (24), (25), and (26), which consisted of three 'context' sentences and a 'target' sentence with a definite reference (the definite noun phrase and its previous mention are in italics):

- (24) *Referent mentioned in Sentence 1:* In one corner of the room was *an upholstered chair*. A broadloom rug in rose and purple colours covered the floor. Dim light from a small brass lamp cast shadows on the walls. *The chair* appeared to be an antique.
- (25) *Referent mentioned in Sentence 2:* A broadloom rug in rose and purple colours covered the floor. In one corner of the room was *an upholstered chair*. Dim light from a small brass lamp cast shadows on the walls. *The chair* appeared to be an antique.

- (26) *Referent mentioned in Sentence 3.* A broadloom rug in rose and purple colours covered the floor. Dim light from a small brass lamp cast shadows on the walls. In one corner of the room was *an upholstered chair*. *The chair* appeared to be an antique.

The target sentence in all three is *The chair appeared to be an antique*, but the referent of *the chair* is mentioned in the first context sentence in (24), the second in (25), and the third in (26). If people have to identify the referent of *the chair* before they are willing to say they understand the target sentence, they should take longer the less accessible that referent is in memory—roughly the farther back that referent was mentioned.

This prediction was confirmed. People were presented with typed versions of three context sentences like (24), (25), or (26). When they felt they understood them, they pressed a button that replaced the context sentences with the target sentence. When they felt they understood that, they pressed the button again and that ended the trial. The average times people spent comprehending the target sentence in sequences like (24), (25), and (26) are shown in Table 9.2. Also listed are the times people took on similar target sentences in which the definite noun phrase was *he*, *she*, *it*, or *them*. As these times show, people were reliably faster on the target sentence—about 300 msec faster—when the referent had been mentioned one sentence back than when it had been mentioned two or three sentences back. It was as if people searched through the entities in short-term memory—those in the last sentence or clause—earlier or more efficiently than they searched through entities not in short-term memory—those in earlier sentences or clauses. In less homogeneous sequences than ours, thematically prominent entities from earlier sentences would probably be in short-term memory too (Chafe, 1974). In any case, the more remotely in memory people had to search for the intended referent, the longer they took to comprehend the sentence, and this is consistent with the proposed model.

Identifying referents should take especially long when listeners have to add bridging assumptions. This prediction was tested in a series of experiments by

Table 9.2 Mean comprehension latencies (milliseconds) for target sentences in one of three contexts. (Referent mentioned in sentence 1, 2, or 3)

Type of noun phrase	Context sentence mentioning referent		
	1	2	3
Nouns	2174	2166	1802
Pronouns	2280	2133	1847
Means	2227	2150	1825

Haviland and myself (Haviland and Clark, 1974). In one study we composed sequences of context and target sentences like the following:

- (27) Esther got some beer out of the car. The beer was warm.
 (28) Esther got some picnic supplies out of the car. The beer was warm.

In (27), the referent for *the beer* is directly mentioned in the context sentence, but in (28), it is not. In (28), listeners have to add the bridging assumption in (28a):

- (28)a. The picnic supplies mentioned include some beer.

If it takes listeners time to form this bridging assumption, they should take longer to comprehend *The beer was warm* in (28) than in (27).

Indeed, they do. In a procedure like the one described earlier, people were timed as they read and comprehended the context sentence and then read and comprehended the target sentence. They took about 200 msec longer, on the average, to comprehend the target sentence in sequences like (28) than they took in sequences like (27). Adding bridging assumptions, then, appears to take time and effort despite the speed and ease with which it seems to be accomplished.

Identifying referents for definite noun phrases, therefore, is a highly inferential activity. Listeners try to infer what is being referred to, and to do that they assume that the speaker has a specific referent in mind that he is confident they will be able to identify uniquely. Only on this assumption can they search their memories for referents they can test for specificity, intendedness, and uniqueness, or can they add bridging assumptions containing a referent that will pass these tests. In this brief section I have examined only some consequences of this view. My main purpose has been to demonstrate that this process is truly inferential, that it depends on tacit agreements, and that it has consequences for theories of comprehension.

SHORTHAND EXPRESSIONS

It is the common view that nouns correspond directly to categories perceived in the world—*neighbourhood* and *Picasso* correspond directly to what people think of as neighbourhoods and Picasso—and when these nouns appear in sentences they are meant to pick out the corresponding world categories and nothing more. In reality, nouns rarely if ever work this way—nor do other more complicated nominal expressions. Instead, they generally refer only obliquely, or indirectly, to the world categories they were meant to pick out. This point can be illustrated with a type of nominal I will call the shorthand expression.

Shorthand expressions are expressions that are rather obviously shortened

versions of longer 'parent' expressions, as shown in these examples:

Parent Expression

three works of art by Picasso
 some Ajax soap
 a Ponderosa pine tree
 six cashmere sweaters
 two shorthand expressions

Shorthand Expression

three Picassos
 some Ajax
 a Ponderosa
 six cashmeres
 two shorthands

Many of these shorthands are so transparent in meaning that they are hardly thought of as shorthands at all. Yet they clearly are. *Picasso*, for example, is a proper noun, and proper nouns cannot themselves take plurals. When *Picasso* does take such an ending, it is short for a parent expression like *people named Picasso*, *people with the characteristics of Picasso*, or *works of art by Picasso*. Similar arguments hold for the other shorthands.

Every shorthand has a multitude of conceivable parent expressions. How, then, do listeners work out which one is intended? Faced with *three Picassos*, how do they decide that it means 'three works of art by Picasso', and not 'three people named Picasso', 'three works of art with the characteristics of one by Picasso', or any of the other conceivable interpretations? There is nothing in the shorthands themselves to say. Once again, listeners have a problem to solve: find the world category the speaker intended his shorthand to pick out. In most instances the problem is solved quickly and easily. Nevertheless, the solution typically requires intricate reasoning.

The interpretation of shorthands rests on a tacit agreement that goes something like this:

The shorthand agreement: The speaker agrees to use a shorthand expression to denote a category (1) that is somehow connected with the explicit content of that expression, (2) that is specific, (3) that is coherent, and (4) that the speaker is confident his listeners can figure out uniquely on the basis of this information.

In illustration, consider the shorthand *William Blake*, as in *I just found a William Blake*. First, the category it denotes must have something to do with William Blake. For example, it might be his paintings, his books of poetry, his poems, quotations by him, books about him, or people with his name. Second, the category must be specific. *One William Blake* cannot mean 'either one painting by William Blake or one book of poetry by William Blake, whichever you happen to think of'. It can mean only one of these at a time. Third, the category must be coherent. *Three William Blakes* cannot pick out two paintings and one book of poetry by William Blake. Such a category lacks coherence. Unfortunately, I cannot yet characterize coherence any more precisely, but

examples like this show that such a property must exist. And fourth, the category has to be one the speaker is confident his listeners can compute—and compute uniquely. If the speaker thinks his listeners do not know that Blake painted pictures, he cannot expect them to interpret *William Blake* as ‘painting by William Blake’ without other hints, as in *I just found a William Blake for my collection of nineteenth-century paintings*. But if he thinks they *do* know Blake painted pictures, he cannot expect them to interpret *William Blake* uniquely as a ‘book of poetry by William Blake’ without hints that eliminate the painting interpretation, as in *I just found a William Blake for my poetry collection*.

Indeed, it is this fourth criterion that demands the subtlest inferences. Consider our implicit reasoning for *I just saw three William Blakes at the art gallery*. First, as objects of *saw*, the entities denoted must be concrete, and that rules out non-concrete objects. Second, as entities seen in an art gallery, we can plausibly rule out all but ‘paintings by William Blake’, ‘portraits of William Blake’, and the like. If the speaker had meant something else, he would have been more specific, since with art galleries we naturally think of art works and the speaker knows that. Third, of the art works connected with Blake, the speaker must have meant Blake’s own paintings. Blake is famous for these, and there are no other well-known art works connected with his name. More to the point, we are confident that the speaker believes we know that, and we are confident that he expects us to use this fact in our reasoning. It is this last assessment that enables us finally to rule out all choices but one.

Two cautions. First, although three *Picassos* may be thought of linguistically as a shorthand for *three works of art by Picasso*, it would be a mistake to assume that listeners reconstruct the parent expression in comprehending it. The model I have proposed assumes instead that listeners search directly for the conceptual category it was meant to pick out. Although this category is described by the parent expression—at least roughly—listeners do not need the parent expression to get at it. Second, listeners do not look for the intended world category itself; they search for its mental representation, however that is specified. Nevertheless, I will continue to speak loosely as if they searched for the world category itself. This is only a figure of speech.

Empirical implications

Although listeners probably never reason as explicitly as these examples suggest, they still cannot arrive at the right interpretations without collecting and weighing evidence as if they did. Somehow they must consult general knowledge and check candidate interpretations against the criteria of specificity, coherence, and unique computability. This view of shorthand comprehension, though formulated less precisely than the views for indirect requests and definite reference, has important empirical implications. I will take up four of them.

The first is that a shorthand will be rejected as unacceptable when it does not satisfy the shorthand agreement. For example, *I just bought five Abraham Lincolns*, without further information, makes no sense, and people reject it, saying, 'You can't say that'. According to the present view they reject it because they cannot find a specific coherent category connected with Lincoln that the speaker could have expected them to identify uniquely. The sentence has so many conceivable interpretations that they are at sea. However, when it is finished out with *for my collection of photographic portraits*, it suddenly becomes acceptable. Now people are able to find a specific coherent category they could have been expected to pick out uniquely.

The trees always vote Republican, without auxiliary information, would be rejected even more swiftly and adamantly than the last example. Unlike *three Abraham Lincolns*, *the trees* contains nothing to signal it to be a shorthand—it is a perfectly good plural—and trees do not vote. Even if it were taken as a shorthand, it is hard to think of sensible things it could pick out. And even if one could do that, it would be impossible to decide on the intended one uniquely. The sentence seems plainly unacceptable. Yet imagine a town divided into voting blocks by area, one in the river bottom, one on the bank, one up in the trees, and one in the meadowland beyond. With that it makes sense to say *The river bottom, the meadow, and the bank generally vote Democratic, but the trees always vote Republican*. *The trees* would be interpreted as 'a majority of the people voting in past elections from the area in the trees'.

Shorthands, therefore, are neither acceptable nor unacceptable in isolation. They are interpreted in relation to certain background information and are acceptable or unacceptable only in relation to that process. The lesson here may be general. To account for the semantic acceptability of a linguistic expression one may always have to refer to the process by which it is interpreted.

This lesson leads directly to the second implication: the interpretation of a shorthand changes with the background information as listeners assess what the speaker knows and why he said what he said. Consider *three William Blakes*, which, as we have seen, has a wealth of conceivable interpretations. There are four broad types of background, or contextual, information that listeners might consult in interpreting it:

1. *Local syntactic constraints*. Compare *I bought* versus *I met three William Blakes*. *Bought* and *met* require the three William Blakes, whatever they are, to be buyable and meetable, respectively.

2. *Global sentential constraints*. Compare *I bought three William Blakes at the art gallery* versus *at the rare book dealer's*. The entities buyable at art galleries and rare book dealers are quite different.

3. *Discourse constraints*. Compare the interpretation of *I bought three William Blakes* in a conversation about paintings versus poetry.

4. *Conversational participant constraints*. Compare *I bought three William*

Blakes spoken by a known collector of poets' letters versus a known collector of first editions.

It is because speakers can usually count on listeners making subtle judgments like these—especially types 3 and 4—that they can use shorthands in a potentially endless number of ways. But it follows that listeners must make just as subtle judgments in interpreting them. Unfortunately, too little is known about how listeners use these constraints.

The third implication is that for a shorthand to pick out the intended category of entities uniquely, it must retain the discriminating parts of its parent expression. Take *some Ajax soap*. In most contexts *Ajax* would be more discriminating than *soap*, and so it would be shortened to *some Ajax*, not *some soap*. Given *some Ajax*, listeners can uniquely pick out soap—Ajax is known for making soap—but given *some soap* they cannot uniquely pick out Ajax—there are many types of soap other than Ajax. Similarly, for the parent expressions given earlier, one would usually say *three Picassos*, not *three works of art*; *a Ponderosa*, not *a pine or a tree*; *six cashmeres*, not *six sweaters*; and *two shorthands*, not *two expressions*.

But the part of the parent expression that discriminates best in one context may not do so in another. Consider *three bottles of Coors beer*. In an opera house bar that sold many kinds of beer all by the bottle, one would order *three Coors*. If the only beer sold was Coors, one would order *three beers*. At a beer stand monopolized by Coors but selling it both by the bottle and by the glass, one would order *three bottles*. But if it sold the beer only in bottles, one would say simply *Three, please*. In the right circumstances all the other parent expressions given earlier can also be shortened in alternative ways.

The part of the parent expression that remains after shortening need not be a noun, as the examples so far might suggest. Indeed, it is hard to find a type of linguistic unit that cannot serve this function, as long as it is pronounceable. Here are examples of several other types of units, each set so that its interpretation is transparent:

- (29) *Adjectives*. A team wearing *yellow* trunks played a team wearing *blue* ones. The yellows beat the blues.
- (30) *Prepositions*. There were so many people driving *to* the supermarket trying to get by those driving *from* it that the *to*'s got entangled with the *from*'s.
- (31) *Conjunctions*. There are people who drink spirits either before *or* after eating, and others who do so both before *and* after. The *or*'s are more sensible than the *and*'s.
- (32) *Prefixes*. Some mothers in the study were interviewed *prenatally* and others *postnatally*. The *pre*'s gave much the same answers as the *post*'s.

- (33) *Auxiliaries*. Twenty people indicated they *would* go to the party, while another ten indicated they *might*. In the end, the *might*'s outnumbered the *would*'s.

In English many shorthands that fall into these categories have already become conventionalized and frozen. For example, *whites* and *blacks* are used for racial groups, *ins* and *outs* for political groups, *ifs* for conditions placed on actions, *pros* and *cons* for people for and against an idea, *haves* and *have nots* for rich and poor people, and *whys* and *hows* for explanations and means. Yet in all these shorthands—the fresh as well as the frozen—the part of the parent expression that is retained is the most discriminating part—typically the unit that forms a contrast with a related expression.

When what remains is a solitary noun, we are faced with an interesting question. Such a noun, in a particular context, can logically be considered one of two things: a full expression, or a shorthand for a fuller expression. In *Robert irritated his brother*, for example, *Robert* may be either a noun phrase complete in itself or a shorthand for *what Robert did*. But can these two cases be distinguished? My present conjecture is that they cannot. Virtually all nouns are intended to pick out categories that are more specific than the categories they literally denote. And if there are nouns that pick out exactly what they say, they are not distinguishable from the rest by any structural characteristics. My assumption is that a noun is rarely what it appears to be, so one must always look for the category it was meant to pick out in that context.

The final implication is concerned with memory: people should give evidence of having picked out a specific world category for each shorthand. Such evidence is available in a study by Barclay, Bransford, Franks, McCarrell, and Nitsch (1974) on memory for sentences like (34) and (35):

- (34) The man tuned the piano.
 (35) The man lifted the piano.

People listened to a list of such sentences and later, with phrases like *something with a nice sound* or *something heavy* as prompts for *the piano*, tried to recall each sentence. Any one person received only one sentence from each of 10 matched pairs of sentences like (34) and (35) and each person was given only one of the two possible prompts for that pair.

What prompts should work best? In (34), *the piano* is short for 'the musical apparatus within the piano', so *something with a nice sound* should be more effective as a prompt than *something heavy*. But in (35), *the piano* is short for 'the physical case of the piano', so in this case *something heavy* should be more effective. This is precisely the pattern that emerged. Prompts were more effective when they cued the specific coherent category the shorthand was intended to pick out. What people remember of a shorthand, therefore, includes its intended interpretation—the real world category it was intended to lead the listener to.

The study of shorthands is still in its infancy—there are far more questions than there are answers. How are world categories, as I have called them, represented in the mind? How do listeners search for the one that is intended? How do they bring the relevant evidence to bear in their decisions? My aim here has been to show how shorthands challenge traditional views of comprehension. Shorthands cannot be accounted for by assuming that listeners look up word meanings in a mental lexicon and mechanically select the corresponding world categories. Shorthands are indirect, and they require indirect means in their comprehension.

CONCLUSIONS

Intentions, therefore, are central to comprehension. Listeners strive to determine what speakers mean by what they say. This is not easy, since speakers rarely mean what they say. The direct meaning of their utterances is only a clue to what they mean, and listeners have to combine such clues with other information and infer the intended interpretations. So a speaker might assert *This soup needs salt* and expect listeners in that context to infer that it also means 'Please pass the salt'. Or a speaker might use *the victim*—describing something as a 'victim'—and expect listeners in that context to pick out the intended referent on that basis. Or a speaker might use *three William Blakes* and expect listeners in that context to see it was meant to pick out three paintings by William Blake. Since inferring what is meant is required for these three expressions, which range from complete utterances to single words, there is little reason to believe it is not also required for most, if not all, other types of expressions.

In this chapter I have focused on how listeners infer what is meant. I have stressed two points. First, comprehension is essentially problem-solving. Second, the problem-solving is accomplished within the confines of certain tacit agreements about language use.

When listeners comprehend an utterance, they in effect solve an intricate problem: What did the speaker mean by what he said? In solving it they use a variety of pieces of information. Why did the speaker say something here now? Why did he select this expression over another? What general knowledge does he expect his listeners to have—from broad facts about common objects and events to specific facts such as what William Blake is noted for? And what does he think they know about this particular context? These and other pieces of information must be gathered and weighed in coming to the right solution. But as in solving other types of problems, listeners sometimes make mistakes, missing a piece of evidence here and weighing a piece of evidence too heavily there. They do not always arrive at the interpretation they were supposed to arrive at.

This solution process relies heavily on tacit agreements people have with each other about the use of language for communication. The basic agreement is Grice's cooperative principle, but it divides into specific agreements. For

indirect requests, there is an agreement about the felicitous performance of the speech act of requesting. For direct reference, there is an agreement about the use of definite noun phrases to enable listeners to compute intended referents uniquely. And for shorthands, there is an agreement about the use of nominals to pick out specific coherent categories uniquely. These agreements are litmus tests any candidate interpretation must pass to be accepted as the intended interpretation. Yet much remains to be learned about these agreements. How precisely are they incorporated into the comprehension process? How do they fit together in a system of agreements about communication? What role do they play in children's acquisition of language?

As the studies I have cited seem to demonstrate, inferring what is meant takes time and effort. Yet it takes so little time and effort that whatever problem-solving goes on must be quick and efficient. Indeed, listeners appear to use heuristic strategies for solving problems they encounter over and over again. The strategies conform to the tacit agreements, but avoid the cumbersome step-by-step reasoning they seem to require. Elsewhere Haviland and I have discussed a strategy listeners appear to use in identifying referents (see Clark, 1977; Clark and Haviland, 1974, 1977; Haviland and Clark, 1974), but little is known about other strategies. For example, given a shorthand expression, listeners seem able to find the right category quickly and without tedious reasoning. But what heuristics do they use in this search? Strategies of this kind, however, cannot be the complete answer, for they sometimes fail. When this happens, listeners seem able to fall back on first principles and use the tacit agreements to reason through what was meant. The strategies are tricks that work most of the time to make problem-solving simpler. Nevertheless, they are still a type of problem-solving.

Comprehension, in short, calls on people's general capacity to think—to use information and solve problems. Although people develop specialized strategies for comprehension, these are still built on their general ability to solve problems—to set up goals, search in the memory for pertinent information, and decide when the goals have been reached. Indeed, in inferring what is meant, people consider non-linguistic factors that are far removed from the utterance itself, and their skill at solving this problem is sometimes taxed to the limit. Comprehension is a form of thinking that should not be set off from the rest.

ACKNOWLEDGEMENTS

The preparation of this chapter was supported in part by Grant MH-20021 from the National Institutes of Health. I am indebted to Eve V. Clark for her comments.

REFERENCES

- Austin, J. L. *How to do things with words*. Oxford: Oxford University Press, 1962.
- Barclay, J. R., Bransford, J. D., Franks, J. J., McCarrell, N. S., and Nitsch, K. Comprehension and semantic flexibility. *Journal of Verbal Learning and Verbal Behavior*, 1974, 13, 471-481.
- Bennett, J. *Linguistic behaviour*. Cambridge University Press, 1976.
- Bolinger, D. L. *Aspects of language* (2nd ed.). New York: Harcourt, Brace, Jovanovich, 1975.
- Bransford, J. D. and Johnson, M. K. Considerations of some problems of comprehension. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.
- Bransford, J. D. and McCarrell, N. S. A sketch of a cognitive approach to comprehension: Some thoughts about understanding what it means to comprehend. In W. B. Weimer and D. S. Palermo (Eds.), *Cognition and the symbolic processes*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1974.
- Chafe, W. L. Language and consciousness. *Language*, 1974, 50, 111-133.
- Clark, H. H. *Semantics and comprehension*. The Hague: Mouton, 1976.
- Clark, H. H. Inferences in comprehension. In D. LaBerge and S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1977.
- Clark, H. H. and Chase, W. G. On the process of comparing sentences against pictures. *Cognitive Psychology*, 1972, 3, 472-517.
- Clark, H. H. and Clark, E. V. *Psychology and language: An introduction to psycholinguistics*. New York: Harcourt, Brace, Jovanovich, 1977.
- Clark, H. H. and Haviland, S. E. Psychological processes as linguistic explanation. In D. Cohen (Ed.), *Explaining linguistic phenomena*. Washington D.C.: Hemisphere Publishing Corp., 1974.
- Clark, H. H. and Haviland, S. E. Comprehension and the given-new contract. In R. Freedle (Ed.), *Discourse comprehension and production*. Norwood, N.J.: Ablex Publishing, 1977.
- Clark, H. H. and Lucy, P. Understanding what is meant from what is said: A study in conversationally conveyed requests. *Journal of Verbal Learning and Verbal Behavior*, 1975, 14, 56-72.
- Fodor, J. A. *The language of thought*. New York: T. Y. Crowell, 1975.
- Gordon, D. and Lakoff, G. Conversational postulates. In *Papers from the Seventh Regional Meeting, Chicago Linguistic Society*, 1971.
- Grice, H. P. Meaning. *Philosophical Review*, 1957, 64, 377-388.
- Grice, H. P. William James Lectures, Harvard University, 1967. Published in part as 'Logic and conversation'. In P. Cole and J. L. Morgan (Eds.), *Syntax and semantics*. Vol. 3: *Speech acts*. New York: Seminar Press, 1975.
- Haviland, S. E. and Clark, H. H. What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*, 1974, 13, 512-521.
- Jarvella, R. J. and Collas, J. G. Memory for the intentions of sentences. *Memory & Cognition*, 1974, 2, 185-188.
- Keenan, J. M., MacWhinney, B., and Mayhew, D. Pragmatics in memory: A study of natural conversation. *Journal of Verbal Learning and Verbal Behavior*, 1977, 16, 549-560.
- Kintsch, W. *The representation of meaning in memory*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1974.
- Sadock, J. M. Speech act idioms. In *Papers from the Eighth Regional Meeting*. Chicago: Chicago Linguistic Society, 1972.

- Sadock, K. M. *Towards a linguistic theory of speech acts*. New York: Academic Press, 1974.
- Schiffer, S. R. *Meaning*. Oxford: Clarendon Press, 1972.
- Searle, J. R. *Speech acts*. Cambridge: Cambridge University Press, 1969.
- Searle, J. R. Indirect speech acts. In P. Cole and J. L. Morgan (Eds.), *Syntax and semantics*. Vol. 3: *Speech acts*. New York: Seminar Press, 1975.
- Tanenhaus, M. K., Carroll, J. M., and Bever, T. G. Sentence-picture verification models as theories of sentence comprehension: A critique of Carpenter and Just. *Psychological Review*, 1976, **83**, 310-317.