1

Definite reference and mutual knowledge

HERBERT H. CLARK and CATHERINE R.
MARSHALL

Jack thinks
      he does not know
   what he thinks
     Jill thinks
     he does not know
But Jill thinks Jack does know it.

So Jill does not know
      she does not know
   that Jack does not know
   that Jill thinks
   that Jack does know
and Jack does not know he does not know
   that Jill does not know she does not know
   that Jack does not know
     that Jill thinks Jack knows
what Jack thinks he does not know

Jack doesn't know he knows
and he doesn't know
   Jill does not know.

Jill doesn't know she doesn't know,
   and doesn't know
     that Jack doesn't know Jill does not know.
They have no problems.

                               Knots, by R. D. Laing

In speaking and listening people make essential use of a great deal of world knowledge that they "share" with each other. The question is, what kind of "shared" knowledge do they use, and how? Recently, in looking at how people plan definite reference, we came on one answer to this question that made us distinctly uneasy. It seemed to suggest that expressions like *the cold asparagus, the mess I made,* and *that animal* require speakers to check a list of facts or beliefs that is infinitely long. Under the most plausible assumptions about how they would actually check that list, they should take an infinitely long time to decide on each noun phrase. However, if there was anything we were certain

about, it was that noun phrases like these are ordinarily selected in a *finite* amount of time – in a few seconds or less. We were at an impasse. The argument for an infinite amount of processing time seemed impeccable, but so did the evidence against it. What we had was a processing paradox, which for reasons that will become clear later we called the *mutual knowledge paradox.*

Like all paradoxes, of course, this one rests on several critical assumptions, and when these assumptions are weakened in one way or another, the paradox can be resolved in several ways. These different resolutions, however, each have their own consequences, and depending on which one we accept, we are led to rather different models for the production and understanding of speech. It is important to decide, then, which way the mutual knowledge paradox is most plausibly resolved.

But we are interested in this paradox only as a way of getting at the two central questions of this chapter: (a) What type of shared knowledge is needed for language use? and (b) how is that shared knowledge in practice assessed and secured? The area of language in which we will take up these questions is definite reference, but even our interest in definite reference is secondary to our concern with the two questions of mutual knowledge. The way we will proceed, then, is to set out the mutual knowledge paradox, describe two ways of resolving it, and argue that one of them is the more usual resolution. We will then suggest that the answers to these two questions bear directly on current theories of language structure and language use, in particular on the characterization and processing of definite reference.

The mutual knowledge paradox

Imagine that there is a Marx brothers film festival on at the Roxy, with one film showing each night for a week. Against this background consider the following scenario:

> *Version 1.* On Wednesday morning Ann reads the early edition of the newspaper which says that *Monkey Business* is playing that night. Later she sees Bob and asks, "Have you ever seen the movie showing at the Roxy tonight?"

Our interest is in Ann's use of the definite referring expression *the movie showing at the Roxy tonight*, term *t*, by which Ann intends to refer to *Monkey Business*, referent *R*. What does Ann have to assure herself of in order to make this reference felicitously? That is, under what conditions does Ann have good reason to believe that Bob won't get

the wrong referent or have to ask for clarification, as with "Which movie do you mean?" The answer we will develop is that she must be certain that once she has made her reference he and she can establish certain shared knowledge about the identity of that referent. Although not all aspects of this scenario are applicable to all other instances of definite reference, we will take up the more general case later.

An obvious first condition is that Ann herself know that the expression *the movie showing at the Roxy tonight* uniquely describes the movie *Monkey Business* – for example, there aren't two movies showing tonight instead. We will describe this knowledge as "*t* is *R*," that is, "the movie at the Roxy tonight is *Monkey Business.*" So, Ann must be certain that after her reference the following condition will be true:[1]

(1)   Ann knows that *t* is *R*.

But is this enough? Obviously not, for what is missing is even the simplest notion of shared knowledge. Specifically, (1) gives no assurance that on the basis of her reference Bob himself will realize that *the movie at the Roxy tonight* uniquely describes *Monkey Business,* a realization that is surely a sine qua non of a felicitous reference. The way Ann's reference may fail can be illustrated by a variation on our original scenario:

> *Version 2.* On Wednesday morning Ann and Bob read the early edition of the newspaper and discuss the fact that it says that *A Day at the Races* is showing that night at the Roxy. Later, after Bob has left, Ann gets the late edition, which prints a correction, which is that it is *Monkey Business* that is actually showing that night. Later, Ann sees Bob and asks, "Have you ever seen the movie showing at the Roxy tonight?"

Although this version satisfies condition (1), Ann has clearly made her definite reference without the proper assurances. She has no reason to think that Bob will realize that the film she is referring to is *Monkey Business.* He is most likely to think it is *A Day at the Races.* The reason why her reference isn't felicitous is clear. She has not assured herself that after she had made her reference Bob will know that *the movie showing at the Roxy tonight* uniquely describes *Monkey Business.* So Ann must satisfy this condition:

(2)   Ann knows that Bob knows that *t* is *R*.

At first, conditions (1) and (2) may appear to be enough, but it is easy to show that they aren't. Consider this variation:

> *Version 3.* On Wednesday morning Ann and Bob read the early edition of the newspaper, and they discuss the fact that it says that *A*

Day at the Races is showing that night at the Roxy. When the late
edition arrives, Bob reads the movie section, notes that the film has
been corrected to *Monkey Business,* and circles it with his red pen.
Later, Ann picks up the late edition, notes the correction and
recognizes Bob's circle around it. She also realizes that Bob has no
way of knowing that she has seen the late edition. Later that day Ann
sees Bob and asks, "Have you ever seen the movie showing at the
Roxy tonight?"

The scenario satisfies conditions (1) and (2). Ann knows that the movie
is *Monkey Business* and that Bob knows that it is too. But she believes
that he believes that she still thinks it is *A Day at the Races.* He is very
likely to take her reference as one to *A Day at the Races* instead of
Monkey Business. Her reference is infelicitous because she hasn't
satisfied this condition:

(3)   Ann knows that Bob knows that Ann knows that *t* is *R*.

The third condition, however, is still not enough, as we can illustrate
with yet another version of the original scenario:

Version 4. On Wednesday morning Ann and Bob read the early
edition of the newspaper and discuss the fact that it says that *A Day at
the Races* is playing that night at the Roxy. Later, Ann sees the late
edition, notes that the movie has been corrected to *Monkey Business,*
and marks it with her blue pencil. Still later, as Ann watches without
Bob knowing it, he picks up the late edition and sees Ann's pencil
mark. That afternoon, Ann sees Bob and asks, "Have you ever seen
the movie showing at the Roxy tonight?"

This version satisfies conditions (1), (2), and (3). Ann knows that the
movie is *Monkey Business;* she knows that Bob knows it too – she saw
him look at the late edition; and she knows that he knows that she
knows it too – she saw him notice her pencil mark on the correct
movie in the late edition. Yet Ann is still not completely justified in
thinking Bob will know she is referring to *Monkey Business.* If she looks
at the world from his point of view, she should reason like this: "She
knows that the movie is *Monkey Business.* But she thinks that I, Bob,
think it is *A Day at the Races,* and so by her reference, she must think I
will pick out *A Day at the Races.*" But if her reference may get Bob to
pick out *A Day at the Races,* it is infelicitous. So we must add another
condition for Ann to be sure of:

(4)   Ann knows that Bob knows that Ann knows that Bob knows that *t* is
      R.

Can we now stop with the confidence that condition (4) is enough?
Not if we can dream up a scenario that satisfies (1) through (4) but still
doesn't justify a felicitous reference. With a little difficulty, we can:

> *Version 5.* On Wednesday morning Ann and Bob read the early edition of the newspaper and discuss the fact that it says that *A Day at the Races* is playing that night at the Roxy. Later, Bob sees the late edition, notices the correction of the movie to *Monkey Business,* and circles it with his red pen. Later, Ann picks up the newspaper, sees the correction, and recognizes Bob's red pen mark. Bob happens to see her notice the correction and his red pen mark. In the mirror Ann sees Bob watch all this, but realizes that Bob hasn't seen that she has noticed him. Later that day, Ann sees Bob and asks, "Have you ever seen the movie showing at the Roxy tonight?"

Complicated as this scenario is, it is possible to see that Ann should not in good conscience have made this definite reference. Putting herself in Bob's shoes again, she should reason like this: "Ann knows that the movie is *Monkey Business,* and she knows that I know that too. Yet she believes that I think she thinks the movie is *A Day at the Races,* and so by her reference, she should think I will decide she is referring to *A Day at the Races.*" But if her reference gets Bob to pick out *A Day at the Races,* it is infelicitous. So we must add condition (5):

(5)   Ann knows that Bob knows that Ann knows that Bob knows that Ann knows that t is R.

Can we be confident that condition (5) is enough? Indeed, we can be sure that it isn't, no matter how fast we seem to be narrowing in on what Ann must be sure of. What these versions show is that there is a way *in principle* of demonstrating that the last piece of embedded knowledge is insufficient. The method is this: Corresponding to Ann's condition (1) is an analogous condition that Bob must assure himself of if he is to be certain she is referring to *Monkey Business.* The condition is this:

(1')   Bob knows that t is R.

For Ann to be sure that her reference succeeds in bringing about this knowledge, she must put herself in Bob's shoes, reason as he would, and make sure she could identify the intended referent uniquely. What we did in constructing Version 2 was to create a scenario in which (1) held after Ann's definite reference, but Ann couldn't know whether (1') held or not. This led us to add condition (2), *Ann knows that Bob knows that t is R,* the equivalent of *Ann knows that* (1'). But just as Ann needs to make sure her reference will bring about (2), Bob has to come to know (2'):

(2')   Bob knows that Ann knows that t is R.

But then (2′) is something else Ann must make sure her reference will bring about, as we showed in creating Version 3, and this led to condition (3). Corresponding to (3), however, is Bob's (3′), which we used in creating Version 4. In principle, we could use this procedure to construct countermanding versions ad infinitum.

The paradox

This view of what Ann has to be sure will result from her use of *the movie showing at the Roxy tonight* suggests a processing paradox. On the one hand, Ann has an infinity of conditions, like (1) through (5), to assure herself of, and that should take her an infinite amount of time. On the other hand, she is surely able to use *the movie showing at the Roxy tonight* as a definite reference, when the circumstances are right, in a finite amount of time. Hence the paradox.

You might rightly complain, however, that the paradox contains a number of hidden assumptions, one or more of which are probably suspect. We see the underlying assumptions to be roughly these:

> *Assumption I.* Ann ordinarily tries to make definite references that are felicitous – ones for which Bob won't get the wrong referent or have to ask "Which one?"
>
> *Assumption II.* To make such a felicitous definite reference Ann must assure herself of each of the infinity of conditions (1), (2), (3), (4), and so on.
>
> *Assumption III.* Each of the conditions (1), (2), (3), (4), and so on takes a finite (though small) amount of time or capacity to check.
>
> *Assumption IV.* Ann ordinarily makes each definite reference in a finite amount of time, on the order of a few seconds.

Assumption I is simply that Ann always tries to make herself understood. She doesn't just blurt out a definite reference and hope against hope that it will work. She chooses her references deliberately and with care. Assumption II merely restates what we have just argued in Ann's reference to *Monkey Business* – that it appears to require her to check an infinity of conditions. Assumption III states a processing assumption that is common to almost every psychological model for such a process – that an infinite number of mental operations cannot be carried out in a finite amount of time (Sternberg, 1966; Townsend, 1972). And Assumption IV states the obvious empirical observation that when people refer to things, they don't take much time in doing it.

The mutual knowledge paradox can be resolved, therefore, by throwing out one or another of these assumptions. Assumptions III and IV seem impossible to get rid of. At least, doing so would take a

great deal of argument. The burden of the paradox, then, falls on Assumptions I and II. Which one, if not both, should we drop? We will return to this question once we have looked more closely at the Frankenstein monster we have created for "shared" knowledge.

"Shared" knowledge

In common parlance, "shared" knowledge has several definitions. Ask your aunt what it means for the two of you to share knowledge that the mayor is an embezzler, and she would probably say, "It means that you know he is an embezzler, and so do I." If p is the proposition that the mayor is an embezzler, then the first definition of shared knowledge comes out like this:

> A and B share$_1$ knowledge that p =$_{def.}$
> (1)   A knows that p.
> (1′)   B knows that p.

Or your aunt might give a more complicated answer: "It means that both of us know that he is an embezzler, and furthermore, I know that you know he is, and you know that I know he is." This leads to a second definition of shared knowledge:

> A and B share$_2$ knowledge that p =$_{def.}$
> (1)   A knows that p.
> (1′)   B knows that p.
> (2)   A knows that B knows that p.
> (2′)   B knows that A knows that p.

We can define a series of types of "shared" knowledge merely by extending the list of statements. These can be denoted by the appropriate subscript on *share*. Shared$_4$ knowledge contains statements down to (4) and (4′), shared$_n$ knowledge, statements down to (n) and $(n′)$. None of these finite definitions, of course, describes the "shared" knowledge required of Ann and Bob after her reference to *Monkey Business*. For that we need something more.

Mutual knowledge

What is required, apparently, is the technical notion of *mutual knowledge*. It has been defined and exploited by Lewis (1969) and Schiffer (1972) for dealing with close cousins of the problem we have raised here. Mutual knowledge is Schiffer's term, whereas Lewis's term for the same thing is common knowledge. We have chosen Schiffer's term, which seems more transparent and less open to misinterpretation. Schiffer defines mutual knowledge as follows:

A and B mutually know that $p =_{\text{def.}}$
(1)　A knows that p.
(1')　B knows that p.
(2)　A knows that B knows that p.
(2')　B knows that A knows that p.
(3)　A knows that B knows that A knows that p.
(3')　B knows that A knows that B knows that p.
et cetera ad infinitum.

Mutual knowledge is the same as shared$_\infty$ knowledge. With the appropriate changes in the definitions, we can also talk about mutual beliefs, mutual expectations, and other mutually held propositional attitudes.

Harman (1977) notes that the infinity of statements in this definition of mutual knowledge can be represented more succinctly in a single self-referential statement of the following kind:[2]

A and B mutually know that $p =_{\text{def.}}$
(q)　A and B know that p and that q.

Cohen (1978) uses a similar representation. In some ways, this definition captures our intuitions about mutual knowledge even better than Schiffer's definition. A visual metaphor will help. Imagine that the proposition p is that the mayor is an embezzler, which Ann and Bob come to know by viewing a picture of the mayor altering the books in the city treasurer's office – he was caught red-handed by a local newspaper photographer. Now by Harman's definition, it is as if Ann and Bob are viewing not only the picture of the mayor's embezzlement, but also a picture of them looking at this picture. That second picture, of course, shows them looking at both pictures, the second of which shows them looking at both pictures, and so on ad infinitum. This definition seems to capture the kind of omniscience Ann and Bob possess about their knowledge of the mayor's embezzlement.

Yet this definition per se doesn't change what Ann and Bob have to assess. Ann must check whether for the Marx brothers example she and Bob know that t is *Monkey Business*. But she must also check to see whether she and Bob know that q, and q is that she and Bob know that t is *Monkey Business* and know that q'. That is, she must check to see whether she and Bob know that she and Bob know that t is *Monkey Business*, and, for q', whether she and Bob know that she and Bob know that she and Bob know that t is *Monkey Business* and that q'', and so on. So just the fact that mutual knowledge can be captured in a single statement doesn't absolve Ann and Bob from checking each of an infinity of statements. Although the representation *looks* simpler, its assessment isn't necessarily simpler.

The form in which mutual knowledge will be most useful, however,

is slightly different from either of these two definitions. Note that both definitions represent mutual knowledge as an omniscient observer would see it, an observer who can say both what A knows and what B knows. But in our Marx brothers examples, Ann was not omniscient. She needed only half the conditions in Schiffer's definition – those numbered *without* primes. It is easy to see that what she needed is equivalent to this:

A knows that A and B mutually know that p.

The effect of this single recursion is to erase all the primes in Schiffer's definition. This assertion says, for example, that A knows that (1'). With (1') spelled out, it says that A knows that B knows that p. But this is equivalent to (2). All the other primes get obliterated in the same way. So from Ann's vantage point, she must determine that she knows that she and Bob mutually know that p, and from his vantage point, he must determine that *he* knows that he and she mutually know that p.[3] Most of the time, however, we will speak informally of A determining merely that A and B mutually know that p, and of B determining merely that A and B mutually know that p. These might be called *one-sided definitions of mutual knowledge*.

Uses of mutual knowledge

The notion of mutual knowledge was originally devised by Lewis to handle some ordinary problems of coordination raised by Schelling in his book *The Strategy of Conflict* (1960). Take the grandfather of all coordination problems:

You are to meet somebody in New York City. You have not been instructed where to meet; you have no prior understanding with the person on where to meet; and you cannot communicate with each other. You are simply told that you will have to guess where to meet and that he is being told the same thing and that you will just have to try to make your guesses coincide. (Schelling, 1960, p. 56)

According to Lewis, you will want to go where the other person will go, namely, where you expect him to go. But you expect him to go where he will expect you to go. Where is that? Where he will expect you to expect him to go, of course. And so on. In short, the two of you will go where you mutually expect the other to go. Whether your mutual expectations are accurate or not is another matter.

If you repeatedly meet your friend at the same place, Lewis argues, you will eventually firm up your expectations and set up a regularity that can be called a convention. It may become a convention, for

example, that the two of you meet, whenever you are supposed to meet in New York City, at the lost-and-found booth of Grand Central Station. But to do so, the two of you must mutually know, among other things, that both of you will go to that booth and that both of you expect each other to go to that booth. In Lewis's formulation, mutual knowledge is indispensable to the definition of convention. It is also, therefore, indispensable to the definition of language because, as Lewis shows, a language like English is in part a system of such conventions.

An application of mutual knowledge closer to our own examples is found in Schiffer's reformulation of Grice's definition of speaker meaning in natural language. Very briefly, his application goes like this. As Grice (1957, p. 385) defined this meaning, "'S [the speaker] meant something by x' is (roughly) equivalent to 'S intended the utterance of x to produce some effect in an audience by means of the recognition of his intention.'" But this definition will not work, Schiffer shows, unless the speaker and audience mutually know, among other things, the effects particular utterances are intended to produce. Schiffer was forced to this conclusion by a series of counterexamples to Grice's definition devised by Strawson (1964), by Searle (1965), and by Schiffer himself. Strawson's and Searle's counterexamples had led to minor repairs in Grice's definition, but Schiffer's, like ours, showed that it was always possible in principle to devise problematic scenarios for "shared" knowledge with fewer than an infinite number of steps. Schiffer's solution was to incorporate the notion of mutual knowledge directly into the definition of speaker meaning, just as Lewis had incorporated it directly into the definition of convention.

Mutual knowledge, then, is ubiquitous. It is an essential ingredient in convention, in meaning, and in language in general. It isn't surprising that it should be an essential ingredient in definite reference too.

Uses of "shared" knowledge

How have other investigators defined "shared" knowledge? Most haven't. The great majority have avoided the problem by not mentioning any interaction between the speaker and listener (for example, J. Anderson, 1976, 1977, 1978; R. Anderson et al., 1976; Ortony and Anderson, 1977; Schank and Abelson, 1977). Others have avoided the problem by limiting the universe of discourse to precisely what the speaker and listener both know. In Winograd's (1972) understanding program, for instance, the commander of the computer "robot" knows what the robot knows and cannot entertain the possibility that there are things the robot knows that he doesn't know. This has been

characteristic of most models within psychology and artificial intelligence.

Within linguistics and philosophy only a handful of investigators have addressed the problem of "shared" knowledge. Several have discussed shared knowledge in a general way, but without saying which kind of shared knowledge they mean. Karttunen (1977, p. 155), for example, talked about a "conversational context," the set of propositions the speaker and addressee can take for granted at that point in the discourse. Later, Karttunen and Peters (1975) introduced the notion of "common ground" or the "common set of presumptions." This consists of the set of propositions "any rational participant [in an exchange of talk] is rationally justified in taking for granted, for example, by virtue of what has been said in the conversation up to that point, what all the participants are in a position to perceive as true, whatever else they mutually know, assume, etc." (p. 286). Karttunen and Peters did not say whether they meant "mutually know, assume, etc." in the technical sense or not. (See also Hawkins, 1978, and McCawley, 1979.)

On several occasions, investigators have committed themselves to specific kinds of shared knowledge. Clark and Haviland (1977), for example, discussed a processing strategy, the given-new strategy, that appeared to require nothing more than shared$_2$ knowledge. Prince (1978) in proposing the notion of tacit assumptions, took up examples that required various amounts of shared knowledge, but didn't bring in anything more than shared$_3$ knowledge. Kempson (1975) explicitly committed herself to shared$_4$ knowledge in discussing the set of propositions that constitute the speaker and hearer's "shared knowledge – knowledge they believe they share" (p. 167). She specifically listed knowledge statements (1), (2), (3), and (4), and no others. In an early paper, Stalnaker (1977) characterized pragmatic presupposition as equivalent to shared$_4$ knowledge: "A proposition P is a pragmatic presupposition of a speaker in a given context just in case the speaker assumes or believes that P, assumes or believes that his addressee believes that P, and assumes or believes that his addressee recognizes that he is making these assumptions, or has these beliefs" (p. 137).

Finally, a few investigators have been explicit in their use of mutual knowledge. In a later paper, Stalnaker (1978) replaced his earlier shared$_4$ knowledge with the notion of "common ground": "Presuppositions are what is taken by the speaker to be the *common ground* of the participants in the conversation, what is treated as their *common knowledge* or *mutual knowledge* [Stalnaker's emphasis]" (p. 321). Similarly, Nunberg (1977), in accounting for definite reference and other

pragmatic problems, introduced the notion of "normal beliefs," which are based on mutual knowledge. And Cohen (1978), in his computational model of speech acts and reference, made essential use of mutual beliefs too.

As this brief survey shows, at least some investigators have felt the need for a notion of shared knowledge. When they have been specific, they have used notions ranging from shared₂ to mutual knowledge. There have been almost as many names for shared knowledge as investigators: conversational context, common ground, common set of presumptions, shared sets, contextual domain, tacit assumptions, pragmatic presuppositions, normal beliefs, and mutual beliefs. Yet these investigators have not taken up the question that would resolve the mutual knowledge paradox: How is shared knowledge assessed in the process of speaking or understanding? Before turning to this question, however, we must take up definite reference itself.

Definite reference

Although definite reference has begotten a vast literature in linguistics, philosophy, artificial intelligence, and psychology, there is still little consensus about its essentials. In this brief section, we cannot hope to do justice to that literature or provide that consensus. Yet to be able to examine the role mutual knowledge plays in definite reference, we need a model for definite reference, no matter how tentative. For this purpose, we will adopt Hawkins' (1978) model of definite reference and modify it a bit to handle some observations of Nunberg (1977) and to make it more closely resemble a related model of Clark and Clark (1979). The only claim we make for this model is that it is a reasonable first approximation – good enough at least to allow us to examine the role of mutual knowledge.

The location theory of the definite article

In Chapter 3 of his book *Definiteness and Indefiniteness* Hawkins reviews the major nongeneric uses of the definite article *the* and then proposes what he calls the location theory of the definite article. He takes up only some uses of the demonstratives *this* and *that;* he doesn't discuss pronouns or proper names. Although his theory is more restrictive than we desire, it is a place to start.

According to the location theory, the speaker performs three acts in using the definite article:

    a.   He introduces a referent (or referents) to the hearer.
    b.   He instructs the hearer to locate the referent in some shared set of objects.

    c. He refers to the totality of the objects or mass within this set that satisfy the referring expression.

To take an example, imagine that Ann told Bob *Bring me the apples.* By this she introduces to him some referents, namely apples. She instructs him to locate apples in some set of objects that she and he share knowledge about. She then refers to the totality of apples in that set – namely, all of the apples. If Ann had said *Bring me the apple,* the shared set of objects would contain exactly one apple, and by referring to the totality of objects within this set, she would have referred to that apple uniquely.

As (b) makes clear, the referent is to be located in a shared set of objects. Where do these shared sets come from? Hawkins argues that they are based on shared knowledge – he doesn't specify which kind – and are inferred "either from previous discourse or from the situation of utterance" (1978, p. 168). As evidence, Hawkins discusses the eight major uses of the definite article put forth by Christopherson (1939) and Jespersen (1949):

1. *The anaphoric use.* In *I bought a lathe, but the machine didn't work right,* the utterance of *a lathe* sets up a "shared previous discourse set," which can subsequently be identified as the referent of *the machine.*

2. *The visible situation use.* In a situation where a bucket is visible to both the speaker and listener, the speaker can say *Pass me the bucket.* The visible bucket constitutes a shared set of objects, which can then be identified as the referent of *the bucket.*[4]

3. *The immediate situation use.* A speaker can use *Do not feed the pony* even though the pony is not visible so long as its existence can be inferred from the situation. Then it is the inferred pony that constitutes the shared set of objects to which *the pony* refers.

4. *The larger situation use based on specific knowledge.* Bob may know the particular store Ann shops at every day, and so it is a shared set of objects. Ann can then refer to it without further explanation, as in *I'm going to the store.*

5. *The larger situation use based on general knowledge.* Ann and Bob know as a general fact that American towns of a certain size each have one city hall. The city hall of Spearfish, the town they happen to be going through at the time, therefore constitutes a shared set of objects that Ann can refer to, as in *I wonder where the city hall is.*

6. *The associative anaphoric use.* In *A car just went by and the exhaust fumes made me sick,* the car is a "trigger" to the "associate" exhaust fumes, and so with the mention of a car, people have a set of associates, which constitutes a shared set of objects. According to Hawkins,

"speaker and listener share knowledge of the generic relationship between trigger and associate" (1978, p. 125).

7. *The unavailable use.* Take *Bill is amazed by the fact that there is so much life on earth,* in which *the fact that there is so much life on earth* introduces new information unknown to the listener. To account for this apparent counterexample to (b), Hawkins takes a transformational approach, arguing that the sentence is derived from *That there is so much life on earth is a fact which Bill is amazed by.* In this source, *a fact* is now indefinite, and so the location theory can be preserved. The unknown information introduced in *The woman whom Max went out with last night was nasty to him,* which contains a "referent-establishing relative clause," is handled in a similar way, and so is the unknown information in *I don't like the color red,* which contains a nominal modifier.

8. *The unexplanatory modifier use.* In *The first person to sail to America was an Icelander,* the definite noun phrase picks out a unique person, whoever he may be, from the set of people who have sailed to America. This is what Donnellan (1966) has called an *attributive* rather than a *referential* use of the definite noun phrase. It picks out "whatever or whoever fits that description," whereas a referential use is "merely one tool for calling attention to a person or thing," and "any other device for doing the same job, another description or name, would do as well" (p. 285). Attributive uses are not intended to secure the mutual knowledge of the identity of the thing being picked out (although they may), and so they should not be assimilated, as they appear to be by Hawkins, with the referential uses, which *are* intended to secure mutual knowledge of the identity of the referent. Our concern is with referential uses, and so we will not consider the unexplanatory modifier use or any other attributive uses any further.

Modifications of the location theory

Like all other current theories of definite reference, the location theory has its problems. At least two of these problems are critical to our enterprise.

The first problem has to do with a condition Hawkins places on the composition of the shared set of objects in (b): "The hearer must either know or be able to infer that the intended object has the property that is used to refer to it in the descriptive predicate" (1978, p. 168). This condition says that for *the ham sandwich,* the hearer must know or be able to infer that the referent is a ham sandwich. This, however, cannot be correct – at least not without qualification.[5] Nunberg (1977) has pointed out systematic examples in which the referent

does not have to have the property of the descriptive predicate. Imagine a waiter in a restaurant pointing to a ham sandwich and saying to another waiter *The ham sandwich is sitting at table six*. In this utterance, *the ham sandwich* is used to refer to the customer who ordered the sandwich, and that customer is obviously not a ham sandwich. Or imagine Ann pointing at her watch and saying *This watch now costs a hundred dollars,* by which she means "an instance of the type of watch this watch is would now cost a hundred dollars" (her own battered watch no longer having much value). Indeed, as Nunberg shows, deferred reference like this is common. To handle such cases, Nunberg introduces the notion of reference function. This is a function the hearer computes on each occasion to get him from the "designatum" (the ham sandwich or watch) to the intended referent (the customer or kind of watch).

The way we will handle this is to distinguish direct from indirect reference precisely on the analogy of direct and indirect illocutionary force. To begin with illocutionary force, *Do you know the time?* can be said to have a *direct* illocutionary force, "Do you have the knowledge of the time?", by virtue of which a speaker can convey a second *indirect* illocutionary force, "Please tell me the time" (Searle, 1975). In his utterance, the speaker intends to convey both illocutionary forces, although the direct meaning may not be intended to be taken seriously (Clark, 1979), and it may convey the indirect meaning by one or another conventional means (Morgan, 1978). Analogously, *the ham sandwich* can be said to have a *direct* referent, the ham sandwich on the plate in front of the waiter, by virtue of which the waiter can indicate a second *indirect* referent, the man who ordered the sandwich. In uttering *The ham sandwich is sitting at table six,* the waiter intends to refer to both objects – the sandwich and the person – although the thing that he is saying is sitting at table six is always the indirect referent. The relation between the direct and indirect referents is determined by Nunberg's reference function.

The condition Hawkins places on the composition of the shared set of objects, then, doesn't need to be changed, as long as we say he is dealing with *direct* definite reference. That is what we will do. We assume that theories of indirect definite reference will proceed along the lines set out by Nunberg as to what constitutes the intended reference function on any particular occasion. As he demonstrates, discovering those functions will not be an easy matter.

The second problem lies in the chronological order of (a) the time of acquisition by the speaker and listener of their shared knowledge of the required set of objects and (b) the moment of the reference act

itself – the moment when the speaker utters the referring expression. Call these two moments Moment$_{SK}$ and Moment$_{RA}$, respectively (SK for shared knowledge and RA for reference act). Although Hawkins doesn't say so explicitly, he seems to assume that Moment$_{SK}$ must precede Moment$_{RA}$. That is, the speaker can only refer to sets of objects he and his listener *already* share knowledge about. This assumption pervades Hawkins' discussion of the first six uses of the definite article, and it seems to motivate his transformational treatment of the seventh.

Is this assumption correct? Clearly not. It appears possible to find counterexamples to the assumption for all eight uses of the definite article. Take the anaphoric use. Contrary to the assumption, it is easy to get an anaphor before its "antecedent," as in *Before he could steal anything, a burglar who had broken into our house was frightened away*. Or take the visible situation use. Contrary to the assumption, Ann can felicitously ask Bob *Please pass the salt* without his realizing there is any salt around. Indeed, it is her reference that induces him to assume there must be salt in view and to look for it. Or take the larger situation use based on general knowledge. Also contrary to the assumption, Ann can felicitously tell Bob *The fourth root of 81 is the number of sisters I have* without assuming that Bob *already* knows what the fourth root of 81 is. She need only assume that he can readily figure it out. In each of these examples, Moment$_{SK}$ comes *after* Moment$_{RA}$, and the shared knowledge is brought about in part by the reference act itself.

But if Moment$_{RA}$ can precede Moment$_{SK}$, there is less reason to posit transformational sources for the seventh use of the definite article. For an alternative analysis, consider Ann's assertion to Bob *The woman Max went out with last night was nasty to him*, where Ann is introducing the woman referred to for the first time. Referent-establishing relative clauses like this, as Hawkins notes, must be anchored to object sets that are already shared, in this instance Max. Ann could not have said, for example, *The woman some man went out with last night was nasty to him*, because *some man* doesn't provide such an anchor. If this is so, Bob can form an object set for this utterance by very much the same procedure as he would for *I wonder where the city hall is*, the fifth use of the definite article. For Ann and Bob it is general knowledge that men like Max often go out with women, ordinarily one woman on any one night, and so Bob can form the set of objects Ann is referring to, namely, the woman Max went out with last night. The requirement, then, seems to be not that Bob already have a shared set of objects, but that he be able to form one based on general or particular mutual knowledge and on the fact that the reference act occurred.

With these two modifications, we can reformulate the location theory in slightly different terms. Tentatively, we suggest the following convention:

The direct definite reference convention. In making a direct definite reference with term t sincerely, the speaker intends to refer to

1. the totality of objects or mass within a set of objects in one possible world, which set of objects is such that
2. the speaker has good reason to believe
3. that on this occasion the listener can readily infer
4. uniquely
5. mutual knowledge of the identity of that set
6. such that the intended objects or mass in the set fit the descriptive predicates in t, or, if t is a rigid designator, are designated by t.

We will not try to justify this formulation in detail, but a few observations are in order. The main point of the convention is this: For a speaker to refer to a thing, he must be confident that because of his speech act the identity of that thing will become mutually known to him and his listener. It doesn't have to be mutually known beforehand, but of course if it were, the listener's inferences would be all that much easier. Ordinarily, to *become* mutually known, the referent must at least be anchored to something that is already mutually known via an anchor cable that is already mutually known. To understand Ann's *I wonder where the city hall is,* Bob doesn't need to believe that the city hall of that town is mutually known, but merely that he and she mutually know about that town (the anchor) and that they mutually know that towns of that size ordinarily have a single city hall (the anchor cable). In condition (6), we have added the notion of a rigid designator, as defined by Kripke (1972, 1977), to take care of proper nouns. In Kripke's theory, *George Washington,* say, is a rigid designator, because it designates the same thing in all possible worlds. In our convention, to use *George Washington* Ann must have good reason to believe that Bob can figure out who it is that the term rigidly designates.

Heuristics for assessing mutual knowledge

For felicitous reference, the speaker and listener must establish certain kinds of mutual knowledge. Simpler notions of "shared" knowledge will not do – as witness Ann's reference to *Monkey Business.* In the light of Lewis's and Schiffer's arguments, this conclusion isn't terribly surprising. Definite reference is an example par excellence of something

speakers and listeners achieve through coordination, and coordination is ordinarily achieved on the basis of mutual expectations. Moreover, definite reference is governed by conventions, and mutual knowledge is an indispensable part of conventions.

But what about the mutual knowledge paradox? It is unthinkable that speakers and listeners assess mutual knowledge by working serially, statement by statement, through an infinity of statements. As we noted earlier, this paradox rests on two debatable assumptions:

> *Assumption I:* Ann ordinarily tries to make definite references that are felicitous.
>
> *Assumption II:* To make a felicitous definite reference, Ann must assure herself of each of the infinity of statements (1), (2), (3), (4), and so on.

The inevitable conclusion is that one or both of these assumptions must be weakened and the infinite process replaced by finite heuristics.

The obvious thing to weaken first is Assumption I. In ordinary speech Ann may sometimes guess at what Bob knows – perhaps guessing wildly– and turn out expressions of definite reference that are far from felicitous. Much of the time this may not matter because her references may be close enough to succeed anyway. And when they don't go through, Bob will look puzzled, ask for clarification, or show other evidence of misunderstanding, and Ann can reassess what she thinks Bob knows and repair her reference. Indeed, repairs of this kind appear to occur often in spontaneous speech, suggesting that speakers don't always satisfy Assumption I with the precision that our Marx brothers examples might have suggested. Perhaps, then, the felicitous reference is an ideal that in practice is rarely reached.

Yet surely it is an ideal people strive for because they will want to avoid misunderstanding whenever possible. What heuristics will enable them to approach this ideal if not reach it? We will suggest two families of heuristics. The first, which we will call truncation heuristics, results in a permanent weakening of Assumption I. The second family, which we will call copresence heuristics, retains the possibility of felicitous definite reference, as in Assumption I, but solves the problems posed by Assumption II.

Truncation heuristics

The stickler in assessing mutual knowledge statement by statement is that there is an infinity of such statements, and that is too many to check. What if people checked only a few of them – like the first four? The task could then be carried out in a finite, even short, period of

time, and that would resolve the mutual knowledge paradox. But if they did this, they could not be guaranteed a felicitous definite reference on each occasion, and Assumption I would no longer hold. Heuristics of this kind will be called *truncation heuristics.*

What makes these heuristics plausible is that they ought ordinarily to lead to few references that are infelicitous. Imagine that Ann always verifies the statement *Ann knows that Bob knows that Ann knows that Bob knows that t is R,* which is condition (4) for mutual knowledge. On actuarial grounds, if condition (4) holds, it should be highly likely that conditions (5) through infinity hold too. So although errors can occur, they should occur rarely and only in complicated situations.

What makes this a *family* of heuristics is that there are several checking procedures a speaker might use. First, imagine that Ann, in referring to *Monkey Business* with the noun phrase *the movie showing at the Roxy tonight,* checks conditions (1) through (4). This might be called the *progressive checking strategy* because Ann starts at the beginning of the list and works so far down. Where she stops depends on her desire for precision. The more precise she wants to be, the farther down the list she will want to check. Second, imagine that Ann checks condition (4) and no others. This might be called the *selective checking procedure.* Once again, the condition Ann picks out to check depends on her precision. The more precise she wants to be, the farther down the list she will want to enter.[6]

Neither of these procedures guarantees a felicitous definite reference because both lead to something less than full mutual knowledge of the referent. Yet in special circumstances there are heuristics that *can* lead to a felicitous reference – so long as the listener draws the right inferences. These heuristics will be called the *augmented truncation heuristics.*

Consider this variation on Version 4:

> *Version 4a.* On Wednesday morning Ann and Bob read the early edition of the newspaper and discuss the fact that it says that there is a double feature playing that night at the Roxy – *Monkey Business* followed by *A Day at the Races.* Later, Ann sees the late edition, notes that *A Day at the Races* has been canceled, and marks the notice with her blue pencil. Still later, as Ann watches without Bob's awareness, Bob picks up the late edition and sees Ann's pencil mark. That afternoon, Ann sees Bob and asks, "Have you ever seen the movie showing at the Roxy tonight?"

Like Version 4, this scenario satisfies conditions (1), (2), and (3), but it also satisfies (4*):

(4*)   There is no R^* such that Ann believes that Bob believes that Ann believes that Bob believes that t is R^*.

Here R^* is a unique referent that fits the description *the movie showing at the Roxy tonight.* Because the reference is singular and there are actually two movies for which condition (4) holds, there is no R^* that fits this description.

These conditions can be enough for a felicitous reference if Ann can count on Bob drawing the right inferences. She could reason this way: "Bob knows that I know that the movie tonight is *Monkey Business.* But because we discussed the early edition, he believes that I believe he thinks there are two movies showing. I can disabuse him of this belief by using a *singular* definite reference. Because he knows I know that *Monkey Business* is the only movie playing, he will infer that I know that he knows that too – even though he doesn't know I know that. He should be able to infer:

(3') Bob knows that Ann knows that Bob knows that t is R.

But because this is my, Ann's, conclusion, I know or believe the equivalent of (4):

(4) Ann knows that Bob knows that Ann knows that Bob knows that t is R.

Reasoning further, I know that this is something Bob could infer, which gives way to (4'), hence my (5), and his (5'), hence my (6), and so on ad infinitum. Voila! He and I mutually know that t is R." In Version 4, it should be noted, Ann could not have reasoned this way precisely because that version doesn't fulfill condition (4*).

When will augmented truncation strategies work? That depends on the precision Ann wants. Imagine a Version 2a (analogous to Version 4a), in which Ann had seen the late edition of the newspaper canceling *A Day at the Races,* but had no idea whether Bob had seen that notice. So she would fulfill:

(1) Ann knows that t is R.
(2) There is no R^* such that Ann knows that Bob knows that t is R^*.

In this version, although she might be sure that Bob realized she was referring to a single movie – he *could* have thought she made a speech error and intended to say *movies* – she has no reason to think he would be able to figure out which one. In an analogous Version 3a, in which she saw Bob look at the late edition but realized he didn't know she had seen it, she could have some confidence he would figure out which. But what Bob really needs to know is that she knows the movie is *Monkey Business,* as in Version 4a. Higher-order versions should make her even more confident he will draw the right inferences.

What we have described, then, is a constellation of conditions that Ann, with certain auxiliary assumptions, can take as good evidence Bob will pick out the right referent. There are probably other such constellations, but all of those we have considered require at least three or four conditions for a reference to be felicitous.

Difficulties with truncation heuristics

In principle, truncation heuristics seem capable of doing the job. They may even allow for felicitous definite reference. We suspect that they may be used on at least some occasions. Version 4a is not such an implausible scenario for people to handle roughly as we suggested. In fact, for scenarios like Version 4a, we have asked subjects to tell us what is being referred to by expressions analogous to *the movie showing at the Roxy tonight.* These subjects appeared to use procedures very much like the truncation heuristics, especially the augmented truncation heuristics. As the scenarios became more complex, they tended to have more difficulty as this analysis would predict. So these heuristics are possible.

But are they plausible as the way people *normally* assess mutual knowledge in making definite reference? We believe not. Our doubts lie in two areas. First, it isn't easy to deal with reciprocal statements as complicated as condition (4). It is implausible that people ordinarily check these conditions per se. Second, the evidence people need in order to verify these conditions anyway suggests a radically different family of heuristics, namely, the copresence heuristics.

Reciprocal knowledge statements, like condition (4), seem unlikely mental objects for people to assess. Recall that in Version 4 of our Marx brothers scenario, we created a situation in which Ann didn't believe that Bob knew that she knew that he knew that the movie that night was *Monkey Business,* a violation of condition (4). The scenario wasn't easy to understand. The main sticking point was in grasping condition (4) and deciding that it wasn't true. Why is condition (4) so difficult to grasp, and to disconfirm?

There are probably two main reasons. One is that recursive statements about propositional attitudes are themselves difficult to grasp. For example, *John Dean knew that Nixon knew that Haldeman knew that Magruder knew that McCord had burgled O'Brien's office in the Watergate Apartments* describes a pipeline of gossip that is difficult to keep straight. When these statements are also reciprocal, with the pipeline turning back on itself, the difficulty seems to increase with the square of the number of recursions. Parallel to the last example is the follow-

ing *reciprocal* statement: *John Dean knew that Nixon knew that John Dean knew that Nixon knew that McCord had burgled O'Brien's office in the Watergate Apartments.* It isn't just that *utterances* of these sentences are difficult to grasp. Rather, their content appears to be inherently hard to keep track of. Any statement more complex than condition (4) can be obliterated by one glass of decent sherry.

Studies of children suggest that the ability to deal explicitly with reciprocal knowledge develops quite late in childhood. In one study (Miller et al., 1970), children were asked to describe cartoons of people thinking of people thinking of people. These children found reciprocal relations much more difficult to describe than nonreciprocal ones. In addition, no more than half the twelve-year-olds were able to deal with reciprocal relations like condition (2), and fewer than a third were able to deal with reciprocal relations like condition (3). In another study (Barenboim, 1978), it was found that children spontaneously talk very little about other people's thoughts (like condition 2) until age twelve, or about other people's thoughts about other people's thoughts (like condition 3) until age sixteen (see also Flavell et al., 1968). All these studies required rather a lot from children – explicit talk about recursiveness and reciprocity – yet they suggest that recursive reasoning even two levels deep is not easy for children under age twelve. The trouble is that children much younger – six to eight years of age (Maratsos, 1976; Warden, 1976) – appear to use definite reference felicitously, at least much of the time. And even younger children sometimes spontaneously repair definite references to take account of what their listeners know (E. Clark and Andersen, 1979). So although studies of children give us anything but a knockdown argument, they do suggest that the truncation heuristics are not very plausible.

The more basic argument against the truncation heuristics is to be found in what counts as evidence for the truth of conditions (1), (2), (3), and so on. Take condition (3), *Ann knows that Bob knows that Ann knows that t is R.* Obviously, Ann won't have this statement represented per se in memory for any arbitrary *t* and *R.* Ann doesn't go through life creating such statements for every object she or anyone else might potentially refer to. Rather, what she needs is a piece of information from which she can deductively or inductively *infer* condition (3). Imagine, for example, a version of our original scenario in which Ann and Bob look at the late edition's correction to *Monkey Business* together. It would be hard to think of better evidence Ann could appeal to for the truth of condition (3).

Ann's knowledge that she and Bob looked at the correction to-

gether, however, is infinitely more useful than that. It is also about the best evidence Ann could appeal to for the truth of all the rest of the infinity of conditions. That is, with this evidence, Ann can jump immediately to full mutual knowledge. If that is so, why would she ever check conditions one by one – even a truncated list of them? She would be better off making sure of the back-up evidence itself. This is precisely the principle that underlies the next family of heuristics, the copresence heuristics.

Copresence heuristics

What kind of evidence can Ann appeal to in order to verify simultaneously the infinity of conditions? If Ann knew, she could in principle satisfy Assumption I and make definite references that were felicitous. She would resolve the mutual knowledge paradox instead by circumventing Assumption II, which otherwise forces her to verify an infinity of conditions one by one. We will argue that what she generally needs is evidence of *triple copresence* – of certain events in which Ann, Bob, and the target object are copresent, as when Ann, Bob, and the notice about *Monkey Business* were openly present together Wednesday morning. The trick is to say what counts as triple copresence – as being "openly present together" – and to say how this can lead to inferences of mutual knowledge.

When Lewis and Schiffer hit on the notion of mutual knowledge, they each recognized the need for a finite means of handling the infinity of conditions. Their solutions were essentially the same. If A and B make certain assumptions about each other's rationality, they can use certain states of affairs as a basis for *inferring* the infinity of conditions all at once. This solution is elegant, for it satisfies everyone. It fits people's intuitions that they mutually know certain facts, and that they yet arrive at this knowledge simply and easily, as if in one short step.

This solution is best illustrated with an example adapted from Schiffer: Ann and Bob are sitting across a table from each other, and there is a single candle between them. Both are looking at the candle, and both see the other looking at it too. The proposition p is that there is a candle on the table. Consider the scene from Ann's point of view. Clearly, she has direct evidence for the truth of (1):

(1)   Ann knows that p.

But she knows other pertinent information too. First, she has evidence that she and Bob are looking at each other and the candle simulta-

neously. We will call this the *simultaneity assumption*. Second, she assumes that he is not only looking at her and the candle, but also *attending* to them. We will call this the *attention assumption*. Finally, Ann assumes that Bob is normal and if he were in her shoes he would be drawing the same conclusions she is. We will call this the *rationality assumption*.

Now if Bob is attending to the candle and is rational, he has evidence for (1'):

(1') Bob knows that *p*.

This, however, is Ann's conclusion, and so she has evidence for (2):

(2) Ann knows that Bob knows that *p*.

But if Bob is rational, he will be drawing the inference that corresponds to hers – his equivalent of (2) – namely (2'):

(2') Bob knows that Ann knows that *p*.

Once again, this is Ann's conclusion, and so she has evidence for (3):

(3) Ann knows that Bob knows that Ann knows that *p*.

In like fashion, Ann would be justified in iterating this process through the remaining knowledge statements (4) through infinity, and Bob would be justified in doing the same for his.

This method for inferring mutual knowledge can be formalized as follows (adapted from Lewis):

> *Mutual knowledge induction schema.* A and B mutually know that *p* if and only if some state of affairs *G* holds such that:
> 1. A and B have reason to believe that *G* holds.
> 2. *G* indicates to A and B that each has reason to believe that *G* holds.
> 3. *G* indicates to A and B that *p*.

G is called the *basis* for the mutual knowledge that *p*. In the candle example, *G* (for "grounds") is Ann and Bob's evidence of triple copresence and their auxiliary assumptions. Ann and Bob each have reason to believe that *G* holds. These grounds *G* indicate to each of them that the other has reason to believe that they hold. And the grounds *G* indicate to both of them that there is a candle on the table. By the induction schema, Ann and Bob mutually know that there is a candle on the table.

The point of this schema is that Ann and Bob don't have to confirm any of the infinity of conditions in mutual knowledge at all. They need only be confident that they have a proper basis *G*, grounds that satisfy

all three requirements of the induction schema. With these grounds, Ann and Bob tacitly realize, so to speak, that they could confirm the infinity of conditions as far down the list as they wanted to go. Because they could do so in principle, they need not do so in fact. This is what gives the copresence heuristics their power. Once one has found proper grounds for mutual knowledge, that is enough.

Mutual knowledge can then be treated as a single mental entity instead of an infinitely long list of ever more complex mental entities. That is, what Ann would represent to herself is not (1), (2), (3), and so on ad infinitum, but merely this: *Ann and Bob mutually know that p.* This obviously leads to an important savings in memory. Just as it is implausible that Ann ordinarily checks a large number of conditions like (1), (2), and (3), so is it implausible that she ordinarily stores these conditions separately in memory. Whenever she needs one of these conditions, she can generate it by a rule such as this (an adaptation of Harman's definition of mutual knowledge): If A and B mutually know that p, then q, where q is that A and B know that p and that q. On demand Ann can deduce, for example, that if she and Bob mutually know there is a candle on the table, then she knows that Bob knows that she knows there is a candle on the table. So with the mutual knowledge induction schema there is simplification in memory too, and the simpler memory structure makes good intuitive sense.[7]

What do the grounds G for the mutual induction schema look like? In the candle example, Ann's grounds consisted of two parts. The first was her direct visual evidence of triple copresence – that there was a candle on the table and that Ann and Bob were simultaneously looking at each other and at the candle. As an event she experienced, this information is relatively fleeting. The second part was her assumptions about the situation – that Bob was consciously attending, that he was doing so simultaneously with her, and that he was rational. These assumptions are more lasting. Ann can assume that Bob is chronically rational, and that if he appears to be looking alertly at a scene, he is attending to it at that moment. These are assumptions she would make for any event of this kind.

There are other grounds too. Some of them are like the candle example but consist of weaker evidence of triple copresence and stronger auxiliary assumptions. For there is a trade-off between the evidence and assumptions loosely as follows:

$$\text{Evidence} + \text{Assumptions} + \text{Induction Schema} = \text{Mutual Knowledge}$$

Because the induction schema is fixed, the weaker the evidence Ann has at her disposal, the stronger the assumptions she must make in

order to satisfy the induction schema and infer mutual knowledge. Still other grounds don't use triple copresence at all. It is instructive to classify the grounds that are most commonly used.

Varieties of mutual knowledge

Mutual knowledge can be classified in various ways. For our purposes it ought to be classified to show its grounds – its sources in a person's experience – because we are interested in how it is secured in the making of definite reference. One main division is between lasting and temporary kinds of mutual knowledge, and another is between several kinds of temporary mutual knowledge. A third division is between *generic* and *particular* knowledge.

Generic knowledge is knowledge about *kinds* of things (about kinds of objects, states, events, and processes), whereas particular knowledge is knowledge about *individual* or *particular* things (about particular objects, states, events, and processes). What we know about dogs in general (that they are animals, that they are domesticated, that they come in many species, and so on) is generic knowledge. What we know about Rin Tin Tin (that he once lived in Hollywood, that he was in several movies, that he was fed caviar, and so on) is particular knowledge. These two types of knowledge are ordinarily expressed in two different ways. Generic knowledge comes in generic sentences like: *Lions roar; A canary is a bird; Rooms each have a floor, a ceiling, at least three walls, at least one door, and they may have windows, carpets, lights, and so on.* Particular knowledge normally comes in nongeneric sentences that refer to particular things, like: *That lion roared just now; Our canary is yellow;* and *The room I am in now has a floor, a ceiling, four walls, two doors, a skylight, a desk, a bookshelf, and so on.* With definite reference, speakers refer to individuals – things in particular knowledge. Yet in doing so, they often need to draw on generic knowledge too.

Community membership

Even when Ann is not acquainted with Bob, she can assume there are generic and particular things the two of them mutually know. The basic idea is that there are things *everyone* in a community knows and assumes that everyone else in that community knows too.[8] In the broad community of educated Americans, for example, people assume that everyone knows such *generic* things as these: Cars drive on the right; senators have terms of six years and representatives terms of two years; and steak costs more than hamburger. They also assume everyone knows such *particular* things as these: George Washington

was the first president of the United States; Colorado is west of
Pennsylvania; there was a great depression between World Wars I and
II. Once two people establish that they belong to the community of
educated Americans, they can assume that they mutually know all of
these things. We will call this mutual knowledge based on community
membership.

But Ann belongs simultaneously to many communities and sub-
communities, each of which has its own distinct areas of knowledge.
At one and the same time Ann could be a high school graduate, a
nineteenth-century-history buff, a San Francisco Forty-Niner football
fan, a psychiatrist, a Palo Alto home owner, an American, a Califor-
nian, a skier, a speaker of Spanish, and a person of Scottish extraction.
For each of these communities, she will have acquired facts she as-
sumes are nearly universal within that community, and she must keep
straight which facts are universal for which communities. She would
not want to meet another person of Scottish ancestry and assume
mutual knolwedge of Freud's theory of neurosis or the Spanish word
for beautiful.

The trick is to judge community membership, and there are many
ways of doing that. Ann may judge Bob to be an American by his
accent, a Palo Alto home owner by his attendance at a meeting of such
home owners, a nineteenth-century-history buff by his description of
the German revolution of 1848, a psychiatrist by his announcement of
that fact, and a person of Scottish ancestry by his surname MacPher-
son. Not only will Ann use these signs in her judgements, but Bob will
provide them intending her to use them for that purpose. In ordinary
conversation people go to some trouble to establish the communities
of which they are members just so that their definite references will
succeed. An illustration of this point can be found in Schegloff's
(1972) account of how people formulate references to places, as when
giving directions.

Before Ann and Bob can assume mutual knowledge of what is
universally known within a community, they must mutually know that
they both belong to that community. Ann might know, for example,
that Bob and she belong to the Stanford University community. But
unless he comes to know that, to know that she knows that, to know
that she knows that he knows that, and so on, he can misinterpret such
references as *the church, the library,* and *the president.* It is easy to imagine
a series of Marx-brothers-like examples that demonstrate this. Ann
could establish mutual knowledge of their Stanford community mem-
bership by her reference itself, as in *Memorial Church, Meyer Library,*
and *Stanford's president,* but this won't always be possible. Requiring

mutual knowledge of community membership introduces a new problem: How do Ann and Bob initially come to mutually know they belong to the same particular community? We suggest that they use one of the copresence heuristics discussed in the next section.

Mutual knowledge of community membership makes an excellent basis G for the mutual knowledge induction schema. Let us suppose that G is "Ann and Bob mutually know that they are both educated Americans." The induction schema requires three things. By requirement (1), Ann and Bob must have good reason to believe G. Indeed, they do. They mutually know they are both educated Americans, which entails that they mutually know G itself. By requirement (2), G must indicate to Ann and Bob that each has reason to believe that G holds. This requirement is fulfilled in the same way. And by requirement (3), G must indicate to Ann and Bob that, for example, American Independence was declared on July 4, 1776. This holds because they assume that every educated American knows the date of American Independence. By the induction schema, it follows that they *mutually* know that American Independence was declared on July 4, 1776.

It is instructive to spell out the two main assumptions required here for mutual knowledge of proposition p. First, Ann must believe that she and Bob mutually know they belong to a particular community. Let us call this assumption *community comembership*. And second, Ann must believe that everyone in that community knows that particular proposition p. Let us call this assumption *universality of knowledge*. Mutual knowledge of this type, then, has a basis G with two assumptions:

1. Community membership: community comembership, universality of knowledge.

Right away we should note two obvious problems. First, communities are not well defined. At what point should a person be considered an educated American, or a member of the Stanford University community, or a nineteenth-century-history buff? Deciding community membership is not a simple task. And second, the two assumptions may vary in strength or certainty. Ann may be certain Bob is an educated American, but less certain that he is a psychiatrist. This is akin to the first point. And she may be more certain an educated American will know that George Washington was the first president of the United States than that Colorado is southwest of South Dakota. The strength of these two assumptions, of course, will affect how certain Ann is that the definite references she is making are felicitous.

Mutual knowledge based on community membership is generally preserved over long periods of time. Once Ann and Bob mutually know they are educated Americans, they are likely to retain that knowledge for use in reference to all sorts of things. And with a constant source of fresh evidence, that mutual knowledge is continually being renewed. Mutual knowledge of the next three types, in contrast, is ordinarily relevant only for short periods of time. It may be used only once and then dropped. Its most distinguishing characteristic is that it is based on evidence that is in a sense more direct.

Physical copresence

The strongest evidence for mutual knowledge that people are generally prepared to accept is what we will call *physical copresence*. An example par excellence is the scene with Ann, Bob, and the candle. Not only are the three of them physically and openly present together, but Ann, say, can readily assume that Bob is attending to this fact, is doing so at the same time she is, and is rational. The physical, or perceptual, evidence Ann possesses is so strong that her three auxiliary assumptions can be relatively trivial. It is rare that she would have reason to think, contrary to the attention assumption, that Bob was catatonic, hypnotized the right way, or very nearsighted, or, contrary to the rationality assumption, that he was too brain-damaged or too young to possess the mutual knowledge induction schema. So with this evidence, once Ann has assured herself of these minimal assumptions, it is trivial for her to refer to *this candle*. Mutual knowledge of the candle has already been secured, and all she has left to do is make sure its identification is unique.

When the time period of physical copresence is placed with respect to the moment of the reference act itself, we can distinguish three varieties of physical copresence. Imagine that Bob isn't paying attention to the target candle, but it is easily within view. Ann can then say *this candle*, which gets Bob to look at it and complete the physical copresence of him, her, and the candle. This could be called *potential physical copresence*. When Ann and Bob are actually focusing on the candle as she says *this candle*, we have a case of *immediate physical copresence*. And when Ann and Bob have looked together at the candle but have stopped before she says *that candle*, we have an instance of *prior physical copresence*.

On the face of it, these three types of physical copresence differ in how strong they are as evidence. The immediate type is the strongest. The potential type is slightly weaker, for Ann must assume that Bob

can discover the target candle and bring it into view simultaneously with her. Let us call this the *locatability assumption*. The prior case is also weaker, for Ann must assume Bob can recall the earlier copresence of him, her, and the candle. Let us call this the *recallability assumption*. If Ann is to use evidence of physical copresence to secure the mutual knowledge necessary for her definite reference, she will need the following auxiliary assumptions:

2. Physical copresence
   a. Immediate: simultaneity, attention, rationality
   b. Potential: simultaneity, attention, rationality, locatability
   c. Prior: simultaneity, attention, rationality, recallability

(Simultaneity, attention, and rationality refer to the assumptions we described earlier.) So far so good. The stronger the evidence, the fewer auxiliary assumptions are needed here.

Linguistic copresence

Many things that are referred to have only been mentioned in conversation. Imagine Ann saying to Bob *I bought a candle yesterday.* By uttering *a candle,* she posits for Bob the existence of a particular candle. If Bob hears and understands her correctly, he will come to know about the candle's existence at the same time as she posits it. It is as if Ann places the candle on the stage in front of the two of them so that it is physically copresent. The two of them can be said to be in the *linguistic copresence* of the candle. Ann can then make a definite reference to the candle, as in *The candle cost me plenty.* [9]

The world in which a thing is claimed to exist can be real or imaginary, past, present, or future. *A deer and a unicorn were grazing beside a stream when the unicorn complimented the deer on his beautiful extra horn.* These two beasts live in an imaginary world, on an imaginary stage, which is quite enough for their linguistic copresence with the speaker and listener. (The question of worlds is too complicated to consider further here, but see McCawley, 1979, and Prince, 1978.)

Unlike physical copresence, linguistic copresence can never be "immediate," that is, simultaneous with the definite reference for which it is used. *A candle* cannot be spoken at the same time as *the candle.* It must come either before, as in *I bought a candle, but the candle was broken,* or afterward if *the candle* is pronominalized, as in *Because it was broken, I returned a candle I had just bought to the store.* In parallel with physical copresence, these two cases can be called prior and potential linguistic copresence, respectively.

To refer to an object that is linguistically copresent, Ann need not use the same term as was used with the potential or prior mention of it. Because a lathe is a machine, and also an inanimate thing, she could say *I bought a lathe, but the machine/it didn't work right.* Note that because not all machines are lathes, it would ordinarily be odd to say, with the same intended interpretation, *I bought a machine, but the lathe didn't work right.*

What auxiliary assumptions are needed for linguistic copresence? To begin with, there are the assumptions of simultaneity, attention, and rationality. Ann and Bob must be attending to Ann's utterance of *a candle* simultaneously, and both must be rational. There is also a complex assumption we will call *understandability.* Ann must assume that Bob will penetrate her indefinite reference, *a candle,* and understand that she is sincerely positing the candle's existence in some world. And as before, prior linguistic copresence requires the assumption of recallability, and potential linguistic copresence the assumption of locatability. Recalling and locating linguistic objects, however, may not be the same as recalling and locating physical objects; so these two assumptions may be either stronger or weaker than those for physical copresence. Putting them all together, we have:

3. Linguistic copresence
    a. Potential: simultaneity, attention, rationality, locatability, understandability
    b. Prior: simultaneity, attention, rationality, recallability, understandability

Fairly clearly, linguistic copresence is ordinarily weaker evidence for mutual knowledge than physical copresence. Whereas seeing is believing, hearing about something requires more – the extra understandability assumption. Both types of copresence are difficult to compare with mutual knowledge based on community membership, whose auxiliary assumptions are so different.

Mixtures

Very often mutual knowledge is established by a combination of physical or linguistic copresence and mutual knowledge based on community membership. Imagine Ann saying to Bob *I bought a candle yesterday, but the wick had broken off.* In uttering *a candle,* Ann establishes the linguistic copresence of him, her, and the candle, but not of him, her, and the wick. To refer to the wick she has to assume that when Bob accepts the existence of the candle, he also accepts the existence of the

wick. He and she mutually know that they belong to the community of educated people for whom it is universally known that candles have wicks. By referring to *the wick,* she can therefore secure mutual knowledge of the identity of the wick that belongs to this particular candle. Ann's use of *a candle,* then, establishes what we will call the *indirect linguistic copresence* of her, Bob, and the wick.

Indirect copresence of this kind may be based on a less certain association than that between candles and wicks (see H. Clark, 1977, 1978; Clark and Haviland, 1977; Haviland and Clark, 1974; and others). For example, a candle has only a *likelihood* of having a wrapper associated with it and only a *low possibility* of being made of bayberries. Yet that is enough to allow Ann to establish their mutual knowledge with her references in *I bought a candle yesterday, but the wrapper was torn* and in *I bought a candle yesterday, and the bayberry smelled great.* Indirect copresence can be very indirect indeed.

There can also be indirect *physical* copresence. A physically present candle, for example, may have a price, which is then indirectly present too. When Ann and Bob are looking at a candle, Ann says *The price was $3,* referring to the candle's price that is indirectly copresent and thereby establishing mutual knowledge of its identity.

Both types of indirect copresence require mutual knowledge based on community membership. That knowledge may be generic, as with candles having wicks, wrappers, bayberries, and prices, but it may also be particular. Imagine that Ann and Bob belong to a small community in which it is universally known that Charlie has a broken left leg. That broken leg is then indirectly copresent with the mention of Charlie. Ann could say to Bob *I saw Charlie yesterday, and the leg is getting better.*

What assumptions are required for inducing mutual knowledge from indirect copresence? If we think of the copresence of the wick as parasitic on the copresence of the candle, then there are first the assumptions of physical or linguistic copresence, whichever is the parasite's host. There is next an assumption we will call *associativity.* It must be mutually known in the community that the parasite is certainly, probably, or possibly a particular part of, or in a particular role with, the host. The two major types of indirect copresence, then, require these assumptions (where parentheses enclose assumptions that are optionally needed depending on the subtype of the host):

4. Indirect copresence
   a. Physical: simultaneity, attention, rationality, (locatability or recallability), associativity
   b. Linguistic: simultaneity, attention, rationality, (locatability or recallability), understandability, associativity

As this listing shows, indirect copresence is always weaker than direct copresence with the parasite's host. The four major types of mutual knowledge are summarized in Table 1.

Types of reference

Traditional linguistic theories tell us that definite reference comes in different kinds. But if definite reference secures mutual knowledge of the identity of R, and if this mutual knowledge is ordinarily inferred from states of affairs G, then definite reference should be classifiable by these grounds G. We will argue that the traditional classifications are indeed based on these grounds G. This argument is important for two reasons. It is indirect evidence that copresence heuristics are used in making definite reference. And it suggests that definite reference cannot be fully explained without bringing them in.

Deixis, anaphora, and proper names

The three basic types of definite reference are deixis, anaphora, and proper names (Lyons, 1977). Deictic expressions are used to point to things in the nonlinguistic situation. In Ann's *I want that, I* refers to the speaker Ann, and *that* refers to the object she is pointing at. Anaphoric expressions are used to refer to things introduced into the conversation itself. In Ann's *I bought a candle, but the thing was broken, the thing* refers to the candle introduced by Ann's utterance of *a candle*. Deixis is often construed to cover anaphora too, but we will stick with its narrower sense. Contrasting with both deixis and anaphora are proper nouns, as in *George Washington had a knotty mouth*. In Kripke's (1972, 1977) proposal, each proper noun rigidly designates the same individual regardless of context.

With this classification, the fit between definite reference and mutual knowledge seems clear. Deixis corresponds to physical copresence; anaphora corresponds to linguistic copresence; and proper names correspond to community membership. The fit could hardly be more obvious. Yet deixis, anaphora, and proper names are categories that are primarily based on functional characteristics. It is worthwhile to look more closely at a few of their structural properties.

The prototypical deictic expressions are demonstratives, as when the speaker gestures at something and says *that,* or *that woman.* These gestures are used to establish immediate physical copresence. They make certain that the speaker and listener come to look at the same object simultaneously. As Hawkins (1978, p. 111) points out, *that* can

Table 1. *Four types of mutual knowledge and their auxiliary assumptions*

Basis for mutual knowledge	Auxiliary assumptions
1. Community membership	Community comembership, universality of knowledge
2. Physical copresence	
a. Immediate	Simultaneity, attention, rationality
b. Potential	Simultaneity, attention, rationality, locatability
c. Prior	Simultaneity, attention, rationality, recallability
3. Linguistic copresence	
a. Potential	Simultaneity, attention, rationality, locatability, understandability
b. Prior	Simultaneity, attention, rationality, recallability, understandability
4. Indirect copresence	
a. Physical	Simultaneity, attention, rationality, (locatability or recallability), associativity
b. Linguistic	Simultaneity, attention, rationality, (locatability or recallability), understandability, associativity

replace a "visible situation use" of the definite article, as in *Look out for the table,* where the table is visible, but not an "immediate situation use," as in *Beware of the dog,* where the dog is somewhere around but not visible. This contrast coincides with our distinction between direct and indirect physical copresence. When there is a candle between Ann and Bob, the candle is physically present, but its price is only indirectly present (unless there is a price tag). Ann can say *That candle is beautiful,* but not *That price is high.* The choice of *that* is governed in part by whether the basis for mutual knowledge is direct or indirect physical copresence.

In anaphora, the prototypical expressions are definite pronouns and definite descriptions, although demonstratives can be used too. In *I met a woman yesterday; the woman/she was a doctor,* the noun phrases *the woman* and *she* are used to refer to a woman already established by linguistic copresence. The type of linguistic copresence is critical. When it has been established in a previous sentence, the speaker can choose either definite descriptions or pronouns, depending on other factors. When it is established in the same sentence as the definite

reference, the choice is highly constrained, as summarized, for example, by Lasnik (1976). In *The woman decided she would operate,* the second reference to the woman must be the pronoun *she.* When there is potential linguistic copresence, it must "command" the definite reference in a technical sense of command. In *She decided that the woman would operate,* it is impossible for *she* to refer to the same person as *the woman.*

Indeed, there appear to be stringent requirements on the basis for mutual knowledge that will allow pronouns. Chafe (1974) has argued that the referents of pronouns must be in the listener's consciousness, "on stage," at that point in the conversation. If so, the conditions on pronouns tie in directly with the assumptions of recallability and locatability. When the referents are recallable, or locatable, within immediate as opposed to long-term memory, the speaker can use a pronoun; otherwise, he cannot.

Demonstratives can be used for anaphora only under special conditions. In *I met a woman yesterday: that woman was a doctor, that* attracts contrastive stress and implies there is a contrasting set of women. It is not used for simple cases of linguistic copresence. And in discourse *this* and *that* are distinguished precisely by the kind of linguistic copresence they require. To refer to something established by prior copresence, one can use either *this* or *that,* but to refer to something yet to be established – potential linguistic copresence – one must use *this.*

Anaphora can be summarized this way. It is prototypically expressed with pronouns or definite descriptions. The expression that is appropriate depends on the type of linguistic copresence: whether it is potential or prior, whether it "commands" the definite reference or not, and whether it is available in immediate or long-term memory, among other things. Anaphora can also be expressed with demonstratives, yet the demonstrative that is appropriate again depends on whether the linguistic copresence is potential or prior. The choice of definite reference, then, is heavily determined by the basis for the mutual knowledge it establishes.

Proper names are the prototypical way of referring to things that are mutually known by community membership. When a particular is widely known in a community, it tends to get a proper name – a rigid designator that doesn't change from one conversation to the next. That is, it is the universally known things within a community that get so named. Note what get proper names: people (*George Washington*), places (*Valley Forge*), and prominent events (*the Revolutionary War*). The few trees, rocks, or animals that get proper names have to be prominent, like *The Great Redwood, Standing Rock,* and *Rin Tin Tin.* There is

probably no grain of sand, glass of water, or ream of typing paper that has ever received a proper name.

Many universally known particulars, however, are referred to with definite descriptions instead of proper names; for example, *the sun, the moon,* and *the snowfall last winter.* Historically, many of these have come to be treated as proper names, as in the change from *the great swamp* to *the Great Swamp,* from *the civil war* to *the Civil War,* and from *the supreme court* to *the Supreme Court.* Sometimes the definite descriptions even become proper nouns, as in *the earth* to *Earth* and *the first world war* to *World War I.* Pronouns and demonstratives apparently cannot be used for reference to things that are mutually known on the basis of community membership, except in rare cases. They require a more direct basis for mutual knowledge.[10]

Eight uses of the definite article

Another classification of definite reference already noted is Christopherson's, Jespersen's, and Hawkins' eight uses for the definite article. Two of these uses are obvious cases of deixis and anaphora. The rest reflect mixtures and fall under our heading of indirect copresence.

The "visible situation use," as in *Pass me the bucket* for a visible bucket, is a clear example of physical copresence, but there are three other uses that are indirect physical copresence. The "immediate situation use," as in *Do not feed the pony* for a nonvisible pony, relies on the physical copresence of a fenced-in yard, supplemented by generic knowledge that such a yard could contain a pony. The "larger situation use based on specific knowledge" relies on the physical copresence of, say, Ann and Bob in a particular situation, with mutual knowledge based on community membership completing the identification of the referent. Ann and Bob mutually know, for example, which store Ann ordinarily goes to in a community; so as long as she and Bob are physically copresent in the neighborhood, she can refer to that store as *I'm going to the store.* The "larger situation use based on general knowledge," as in *I wonder where the city hall is* for a new town, has a similar basis.

The "anaphoric use" is a plain example of linguistic copresence, and the "associative anaphoric use," as in *A car just went by and the exhaust fumes made me sick,* a case of indirect linguistic copresence. Within a community, cars are known to produce exhaust fumes, knowledge that along with the linguistic copresence of the car is enough to secure mutual knowledge of the fumes being referred to.

There are several subtypes of "unavailable use" of the definite article. The first is the "referent establishing relative clauses," as in *The woman Max went out with last night was nasty to him.* These always "relate the new, unknown object [here, the woman] either to other objects in the previous discourse set, or to participants in the speech act, or else they identify entities in the immediate situation of the utterance" (Hawkins, 1978, p. 137). So they are cases of indirect linguistic or physical copresence via mutual knowledge based on community membership. A second subtype is the "associative clause," as in *the beginning of World War II,* another case of indirect copresence because it is mutually known among educated people that wars have beginnings, and so one can indirectly identify the beginning of a mutually known war. A third subtype is the "noun complement," as in *the idea that he is in Caracas* and *the fact that the world is round.* One way to view these is to say there is a set of possible ideas, and a set of possible facts, and these sets are mutually known based on membership in the community of thinking, perceiving humans. Any individual fact or idea can then be identified merely by being specified. The fourth type, "nominal modifiers," as in *the color red,* would work the same way.

Deixis as fundamental

According to many linguists (for example, Lyons, 1975), deixis is the source for all definite reference. In Indo-European languages, the pronouns (like English *he, she, it,* and *they*) and the definite articles (like English *the*) are historically derived from demonstratives (like English *this* and *that*). Thorne (1972, 1974) has argued that the definite article is fundamentally locative – that is, deictic. *The woman* designates not merely a particular woman, but a particular woman in a particular place. All the world's languages appear to have demonstratives and personal pronouns, but many do not have definite articles. In these languages, when a definite reference has to be made absolutely clear, a demonstrative is used, as in *that woman* (Moravcsik, 1969). That is, demonstratives are stretched to cover other nondeictic kinds of definite reference. And in language acquisition, E. Clark (1978; Clark and Sengul, 1978) has argued that deixis is also fundamental. Children refer to things by pointing long before they begin to speak, and their first referring expressions, usually *that, there,* or *look* in English, are almost invariably accompanied by pointing. The weaker forms of definite reference – the pronouns and definite article – are acquired only later. Proper nouns, our incommensurate case, however, come in very early (E. Clark, 1973).

If, as we have argued, physical copresence is the fundamental type of copresence, then it follows that deixis should have primacy in definite reference. The idea is this: Physical copresence is the prototype of what it means for a thing to be mutually known. It is such good evidence that it needs only weak auxiliary assumptions to serve as a basis G in the mutual knowledge induction schema. The other types of copresence each require stronger assumptions, as if they were defective types of copresence in which one or another of the essential conditions of physical copresence hadn't been fulfilled. If physical copresence is primary, then deixis too should be primary. It is significant that there is such a convergence of evidence from historical linguistics, language universals, and language acquisition.

To summarize, when definite reference is divided into types, these correspond to different bases G by which mutual knowledge of the identity of the referent is established. And among these types, deixis appears to be primary. All this evidence is in line with the copresence heuristics – in particular with the use of physical, linguistic, and indirect copresence and of community membership.

Reference repairs

In conversation people often say one thing, repair what they have just said, and then go on (see Clark and Clark, 1977, pp. 260–71). Ann might say *I ran into Ralph – you know, the guy who works in our clinic – the other night at the symphony*. Or she might say *I ran into Ralph the other night at the symphony*, to which Bob would ask *Who is Ralph*, to which Ann would reply *You know, the guy who works in our clinic*. Both types of repairs – self-repairs and other-repairs – are common in everyday speech, although self-repairs predominate (Schegloff et al., 1977).

Repairs of definite reference, what DuBois (1975) has called *reference editing*, give further evidence for the copresence heuristics. The argument is this. One reason speakers repair definite references is to make them more likely to succeed. In our examples, Ann wants to make it more likely Bob will identify the person she was referring to. In making these repairs speakers have two broad options. They can provide more information in the reference itself. This way the basis G on which they and their listeners come to mutually know the identity of the referent will become clearer. This might be called a *horizontal repair*. Or they can strengthen the type of copresence on which their reference is based. This might be called a *vertical repair*. Of these repairs, some should increase the success of a reference, and others should not. If our proposal is correct, those that increase success

ought to be just those that provide stronger types of copresence – direct instead of indirect, physical instead of linguistic, immediate instead of potential. The evidence is that they are.

Horizontal repairs

Most reference repairs are horizontal. They ensure greater success by providing more precise information about the referent without changing the type of basis G on which its identity becomes mutually known. Consider these four cases.

1. *Physical copresence.* Imagine telling a librarian with a gesture, *I want that.* He prompts, *Which one?* You reply, *The book right there on the second shelf.* He prompts again, *I still don't see which one.* You reply, *The green book on the second shelf from the bottom of that bookcase.* These references all rely on potential physical copresence. What changes with each repair is the precision with which the referent is specified. This you accomplish by adding descriptors that refer to other potential physically copresent items – *right there, the second shelf, the bottom, that bookcase.* Each addition, you believe, makes it more likely that the right book will be identified uniquely. Each new piece of information strengthens the basis G on which the identity of the referent can be mutually known. Horizontal repairs of *prior* physical copresence work the same way.

2. *Linguistic copresence.* Imagine this interchange. Ann: *A doctor I met last night introduced me to a lawyer, and she gave me some advice.* Bob: *Who did?* Ann? *The lawyer.* In this repair Ann has disambiguated her reference by providing one more descriptor – that the referent is a lawyer. This descriptor, like the reference itself, is based on prior linguistic copresence. She could not have added *the woman in black* or *the person near the piano* or *the rich one,* which do not make contact with information Ann has provided linguistically, but she could have said *the person the doctor introduced me to* or *not the doctor* or *the one I talked to second.* To be effective, horizontal repairs must add or alter descriptors, not delete them. It wouldn't make sense for Ann to say *A woman I met last night introduced me to her daughter, and the older one, I mean she, gave me good advice.* Cooperative repairs – and that is what we are talking about – must lead to a more precise identification of the referent.

3. *Indirect copresence.* Imagine Ann's report: *I tried to get downtown yesterday but the bus – the one I was riding in – broke down.* The bus is identifiable only on the basis of indirect copresence, and the repair adds other evidence of indirect copresence, namely that Ann *rode* on the bus. Like the previous two types of repairs, the more information

the listener is provided with, the more successful the reference is judged to be.

4. *Community membership.* When references rely on community membership, there are several ways of making horizontal repairs. One is to add more information, as in *I met Nina – Nina Baker,* or as in *I hated the war – the Vietnam War.* Another is to change the community basis for the reference, as in *I like my new colleague – you know, Elizabeth Adams.* Here both the original reference and the repair rely on community membership. What the repair does is change the community from one in which it is universally known who the speaker's colleague is to one in which it is universally known who Elizabeth Adams is. This change in community must strengthen either the certainty it is mutually known that the speaker and listener belong to that community or the certainty that the referent is universally known in that community.

Vertical repairs

The principle of repairs is that they *strengthen* the basis G on which mutual knowledge of the identity of the referent can be inferred. With this principle we can examine vertical repairs, ones that replace one kind of copresence by another, to see if we can order the types of copresence for their strength. If our proposal is correct, the types of copresence should order themselves from strongest to weakest according to the number and kind of auxiliary assumptions they require. Indeed, that is what we will demonstrate.

1. *Physical copresence.* Among types of physical copresence, immediate physical copresence should be the strongest because it requires the fewest auxiliary assumptions – and it is. Ann: *The book over there is mine.* Bob: *Which one?* Ann, picking up a book and showing it to Bob: *This one.* In her repair Ann has moved from potential to immediate physical copresence. If she had moved in the reverse direction, from immediate to potential physical copresence, her repair would have been nonsense. Or imagine Ann: *The book I just showed you is mine.* Bob: *Which book was that?* Ann, picking up a book and showing it to Bob: *This one.* Ann's repair here goes from prior to immediate physical copresence, and it too would be nonsensical in the reverse direction. Recall that the basis G for potential and prior physical copresence requires the auxiliary assumptions of locatability and recallability, respectively. What these and similar repairs show is that some such assumptions are necessary and that G can be strengthened by turning to direct evidence that doesn't need them – namely, immediate physi-

cal copresence. As Searle (1969, p. 88) has argued, the limiting case of referring to something is physically showing it (along with a suitable expression).

Physical copresence is stronger when it is direct than when it is indirect. Ann, still staring at Bob and the candle: *The price was too high.* Bob: *What price?* Ann: *The price of this candle.* And physical copresence is stronger, all other things being equal, than linguistic copresence. Ann: *I was just reading a book on your bookshelf, and it was terrific.* Bob: *What book?* Ann, picking out a book and showing it to Bob: *This one.* Repairs like these, then, are evidence for the auxiliary assumptions of associativity and understandability that we said were required for indirect physical copresence and for linguistic copresence.

2. *Linguistic copresence.* When things are not physically showable, repairs have to be made that move up to the strongest kind of linguistic copresence. Ann: *I bought a candle today; the seal was broken.* Bob: *What seal?* Ann: *The seal on the wrapper around the candle.* This repair moves the evidence up from indirect to direct linguistic copresence. It suggests that the assumption of associativity is not a trivial one. And within linguistic copresence, a repair can be made that strengthens the recallability or locatability of the linguistic copresence. Ann: *I think your idea is excellent.* Bob: *What idea?* Ann: *A moment ago you mentioned going to a movie tonight.* This repair brings back into linguistic copresence an idea Bob had failed to recall.

3. *Community membership.* Community membership cannot be ordered for strength in relation to physical, linguistic, or indirect copresence because its auxiliary assumptions are not comparable with those of the other three types. It can apparently be either stronger or weaker than physical or linguistic copresence, depending on the purpose for the repair. Take this exchange: Ann: *I was just talking to the woman standing right over there* (pointing). Bob: *Who is she?* Ann: *Nina Baker, the artist.* Contrast it with this interchange: Ann: *I was just talking to Nina Baker, the artist.* Bob: *Who is she?* Ann: *The woman standing right over there* (pointing). In the first exchange, the woman's physical appearance was not as significant an identification for Bob as her role in Ann and Bob's community. In the second, it is the other way around, as if Bob knows little about Nina Baker in the community and now at least can identify her physically. These repairs bear out the claim that the community membership has auxiliary assumptions that are not comparable with the others.[11]

The several bases G we proposed earlier – physical, linguistic, and indirect copresence and community membership – are only one way of cutting up the territory. They provide a tidy geography in which

each basis has associated with it a few assumptions, such as simultaneity, recallability, understandability, and community comembership. A more thorough survey of repairs might suggest a different geography with slightly different auxiliary assumptions. Still, such a survey would rely on the logic we have just been using. Every repair that is judged to strengthen a reference should be associated with the elimination or simplification of one or more auxiliary assumptions. Such a survey should lead to a more complete map of the copresence heuristics themselves.

To summarize, repairs of definite reference bear witness to people's use of copresence heuristics. When a speaker makes such a repair, he tacitly reassesses his evidence for mutual knowledge of the identity of the referent, and his repair is an attempt to strengthen that evidence. The way he strengthens it is to try to find fresh evidence that needs weaker or fewer auxiliary assumptions.

Organization of memory

The copresence heuristics, with their voracious appetite, can be satiated only by the right kinds and amounts of factual fodder. How is this fodder organized? What does the storehouse of data the heuristics feed on look like? The arguments we have offered so far suggest a rather different view of memory from those of most current models of understanding and production.

One traditional view of definite reference is that its primary function is to pick out particular individuals – individual objects, states, events, or processes (see Strawson, 1974). What this view has suggested to most investigators is that in processing definite reference people search memory for the particulars actually referred to. They can't, of course, find the particulars themselves, but they can find *referential indices* corresponding to them. Each index is a stand-in, so to speak, for the referent itself. Imagine that Bob's memory contains a set of referential indices for entities represented as $E_1, E_2, E_3, \ldots E_n$, and that E_3 is the referential index for *Monkey Business*. When Ann uses the definite description *the movie showing at the Roxy tonight,* he is supposed to search this list and settle on the intended referent E_3. Although current models of comprehension differ in their specifics, virtually all of them assume this kind of search for the intended referent, including those of Anderson (1976), Clark and Haviland (1977), Kintsch (1974), Kintsch and van Dijk (1978), Rumelhart et al. (1972), Schank and Abelson (1977), and Winograd (1972), to name just a few.

All of these models, however, are incomplete. Bob cannot search

memory for E_3 alone, for that would hardly guarantee that E_3 was mutually known to him and Ann. In most cases, he must search for an *event* that involves not only E_3 but also E_1, Ann, and E_2, him. This event, call it E_4, has to be evidence of their physical, linguistic, or indirect copresence. Or when community membership is concerned, he must search for an individual E_3 that everyone in a community he (E_1) and Ann (E_2) both belong to knows. In none of the models just mentioned does the listener search for such an event or for such a community-wide individual.

Components of memory

Our point can be made with a metaphorical view of memory as a personal archive, or library, in which there are several different kinds of reference books. Most theories of understanding require memory to contain a grammar of English and a dictionary. With these two books, the listener can parse sentences and figure out what they mean. But to handle definite reference, the listener needs more.

What most current models of comprehension add is an elaborate kind of telephone book. In a definite reference like *the man in the red shirt*, Bob is told the name and address of the individual whose referential index he is seeking. All he needs to do is search the telephone book for this name and address, and the book will tell him the right referential index – the telephone number that connects his name and address (the reference) with his physical person (the referent). The telephone book must be a sophisticated one, like the Yellow Pages, in which the names and addresses are organized and cross-classified according to some scheme. But in effect it is a mere listing of descriptions of individuals paired with their referential indices.

Such a telephone book won't do, however, because it doesn't contain the right kind of information. Take Ann's telephone book. For her to be able to make a successful definite reference, the book would have to distinguish those names and addresses she knew Bob knew from the rest, and it would have to make the same distinction for everyone else she might potentially talk to. Although that satisfies condition (2) for mutual knowledge, it doesn't do anything more. Her book would also have to distinguish those names and addresses she knew Bob knew she knew from the rest, satisfying condition (3), and those she knew Bob knew she knew Bob knew from the rest, satisfying condition (4), and so on. Very quickly, her book would grow unmanageably large. The telephone book, in effect, is an embodiment of Assumption II, which is just the assumption we want to circumvent in order to avoid the mutual knowledge paradox.

What the copresence heuristics require instead is a pair of books, a diary and an encyclopedia. Bob's diary is a personal log that keeps an account of everything significant Bob does and experiences. When Ann uses the reference *the man in the red shirt,* Bob must find in memory an individual who fits that description – a man in a red shirt. But he knows that he must search his diary for an entry that gives evidence of the physical, linguistic, or indirect copresence of him, Ann, and that man. That is, he must seek out an *event* that he can use along with certain auxiliary assumptions as the basis *G* for inductively inferring mutual knowledge of the identity of that man. This is far more complicated than searching a telephone book for a number. Every event he searches for involves the referent plus two other individuals, and that takes more specification than the referent alone.

Not all parts of the diary will be equally accessible. The more recent events ought to be more accessible, and there is evidence to suggest that they are. In several studies, people were found to take less time to understand definite references that relied on linguistic copresence the more recently the antecedent event occurred (Carpenter and Just, 1977; Clark and Sengul, 1979; Lesgold et al., 1979). And events that are more significant ought to be more accessible too. However, too little is known to be able to say much more about the organization of the diary. Our point is that such a diary is needed to account for genuine cases of felicitous definite reference.

Bob's second book is an encyclopedia, which he needs for mutual knowledge based on community membership. It will have recorded in it all the generic and particular knowledge Bob believes is universal to each community he belongs to. Instead of being organized in the conventional way – alphabetical by subject matter – it might take this form: Chapter 1 would contain the knowledge every human being is assumed to know, Chapter 2 the additional knowledge every American is assumed to know, Chapter 3 the additional information (over Chapters 1 and 2) that every Californian is assumed to know, and so on. Within each chapter there would be sections on biographical, geographical, historical, and other types of information. And there would be special chapters for the additional specialized knowledge possessed by psychiatrists, by Palo Alto homeowners, and by whatever other communities and subcommunities Bob may happen to belong to. Happily, subject matter and communities tend to go hand in hand – psychiatry is known by psychiatrists, and the rules and regulations for owning homes in Palo Alto by Palo Alto homeowners – and so the encyclopedia doesn't have as complicated an organization as it might first appear.

It is the encyclopedia that Bob consults for references that require

mutual knowledge based on community membership. Imagine that Bob and Ann mutually establish through months of companionship that they both belong to certain communities – those corresponding, say, to Chapters 1 through 8, 11, 15, and 33 in Bob's encyclopedia. When Ann uses the reference *George Washington,* Bob must search just those chapters for an individual with that name. He must also consult those chapters for her references that rely on indirect copresence. When she says *I went to buy a candle but the price was too high,* he will find what is known about candles, determine that each has a price, use this information to create an individual (or rather, its referential index) that corresponds to the price of the candle she mentioned, and identify it as the referent of *the price.* Creating such referential indices via indirect copresence is known to take people longer than merely identifying referential indices that are already present. People understand *The beer was warm* more quickly after *Mary got some beer out of the car,* where the beer is directly copresent, than after *Mary got some picnic supplies out of the car,* where the beer is only indirectly copresent (Haviland and Clark, 1974).

A great deal has been said about the organization of such an encyclopedia. Minsky (1975) has proposed that people have "frames" for what such things as rooms consist of in general and what specific rooms consist of. Schank and Abelson (1977) have made a similar proposal for "scripts" of what people should and actually do do in such activities as going to a restaurant. Rumelhart and Ortony (1977) have proposed "schemata." Yet in none of these proposals is there any consideration for how this knowledge might be compartmentalized according to what information is mutually known by a community or by two individuals, as required for definite reference.

The diary and encyclopedia are not independent of each other. They must be cross-indexed by the individuals they contain – as when someone speaks of George Washington, the Revolutionary War, and 1776 and then refers to them all in *He led the army then.* And certain diary entries will be duplicated in the encyclopedia, as when Bob sees a news item on CBS television and supposes that it is universally known by the community of people who habitually watch CBS television.

Speaker models and listener models

The memory described so far seems entirely too large and unwieldy for everyday use. It seems to go against people's intuition that talk is easy, that getting the right information at the right time is effortless and straightforward. Their intuition is based, we suggest, on the fact

that the diary and the encyclopedia are compartmentalized into useful units. In conversation the units that are pertinent at any time can be prepared for selective access.

Imagine, at a party, turning from talk with an English speaker to talk with a French one. You are likely to feel you are changing gears – as if you are putting away your English dictionary and grammar and pulling out your French ones. Our suggestion is that you make similar shifts whenever you change interlocutors. You prepare yourself selectively to talk to, or listen to, that particular person or group of people. You do this by selecting pertinent parts of the diary and encyclopedia for ready access.

The way a speaker prepares is by accessing his model of the listener, and the listener accesses his model of the speaker. When Ann talks to Bob, she creates in memory a model of what is in Bob's mind – his knowledge, his perceptions, his current thoughts – and she constantly updates it. Bob carries along a similar model of what is in Ann's mind. These models must include the right diary entries and encyclopedia chapters. Ann's model of Bob would contain all those chapters of her encyclopedia that correspond to communities she knows he belongs to. However, she knows she can refer only to individuals in those communities she knows they mutually know they *both* belong to. She may know Bob is a chess addict, but realize he doesn't know she knows. So her model may include Chapters 1 through 11, 16 through 24, 38, and 55, but of those only Chapters 1 through 8, 16 through 18, and 55 are mutually known. Her model of him also contains all those diary entries that involve Bob in some way. It is these she consults when deciding whether she can establish mutual knowledge of the identity of most individuals she wants to refer to. Ann's model of Bob, in short, contains just those parts of her diary and encyclopedia that will be useful for getting him to understand her, whatever she may want to talk about. It will also contain just those parts that will allow her to understand him and all his actions.

The suggestion is that we carry around rather detailed models of people we know, especially of people we know well. If Bob is a close friend of Ann, she may even have a special chapter in her encyclopedia for him, as if he and she form a community of two people. It is hard to underestimate the importance of these models. At a cocktail party, as Ann turns from Bob, her close friend, to Charles, her cousin from out of town, her model of the listener will change radically, and so will the way she refers. Diary entries are particularly important here. If she has just told Bob about her theory of the Marx brothers' success in Hollywood, she cannot immediately expect Charles, who

has not heard what she told Bob, to understand references to things she told Bob. She must keep track – careful track – of what she told each of them. Though her Marx-brothers theory may be uppermost in her mind as she turns to Charles, she has to explain it over again to him if he is to understand her. People who tell someone the same gossip, or joke, or piece of news twice without realizing it are considered impolite or absentminded. They have failed in the social imperative of keeping their models of each particular listener straight.

How do we build these models in the first place? In certain circumstances we can watch our model of a person being erected block by block. One of these is in formal introductions, which are designed to lay the foundations of our model of the other person in the first few seconds and to add onto it prefabricated sections quickly and easily.

Imagine Ann at a party of academics bumping into Ed. "Isn't the weather just great!" she tells Ed. The weather one can always refer to, because it is mutually identifiable by people in the same locale. The convention of always talking about the weather at the beginnings of conversations and in new conversations has an obvious basis in mutual knowledge.

"Yes, it is," replies Ed. "My name is Ed Taylor. I'm a psychiatrist working here at the Palo Alto VA hospital." With this, Ann can add to her model of Ed not only, say, encyclopedia Chapters 1 and 2 – for being human and being American, which she could gather from his reply alone – but also Chapters 3 through 11, 15 and 25, for being a Californian, a Palo Altoan, a psychiatrist, and so on. This is typical of self- and third-party introductions. They allow one to build up great chunks of the model of the other person. They are intended to accomplish just that so that the two now have something to talk about, things they can felicitously refer to.

"How do you do. And I'm Ann Horton, and I work in the psychiatry department at Stanford." With this, Ann has established mutual knowledge of the universal information in these chapters. She was able to refer to Stanford University and its psychiatry department just because she knew Ed was a member of the Palo Alto community and the community of psychiatrists, and so her reference would secure mutual knowledge of the identity of these two places. "What kind of psychiatry do you specialize in?" she might go on, continuing to establish mutual knowledge of larger and larger spheres of experience.

In summary, people's memory must be organized to enable them to get access to evidence they will need to make felicitous references. What that implies is that their memory must contain a diary of significant personal experiences cross-indexed with an encyclopedia or-

ganized both by subject matter and by the communities who possess the knowledge. It also suggests that people have selective access to information that is pertinent to each person they talk to. They have a model of what is in the other person's mind, a model they have built up from previous contact and which they continue to update as they go on talking. It is that model that enables people to make and understand references so quickly and accurately.

Conclusions

Definite reference is one of those phenomena in language that seem so obvious that it is hard to see what there is to explain. We have tried to shatter this illusion by posing the mutual knowledge paradox, which is this: To make or interpret definite references people have to assess certain "shared" knowledge. This knowledge, it turns out, is defined by an infinite number of conditions. How then can people assess this knowledge in a finite amount of time? From the beginning, we knew the paradox was illusory – one or more of its assumptions had to be incorrect. Yet we found it a useful magnifying glass for looking into the processes by which people use and understand definite reference.

The resolution of the paradox we favor for most circumstances is that people assess mutual knowledge by use of the copresence heuristics. They search memory for evidence that they, their listeners, and the object they are referring to have been "openly present together" physically, linguistically, or indirectly. Or they search memory for evidence that the object is universally known within a community they and their listeners mutually know they belong to. With such evidence they can infer mutual knowledge directly by means of an induction schema. There is no need to assess an infinite number of conditions, and the paradox collapses.

The copresence heuristics have important consequences for definite reference. They help determine people's choice of noun phrase for each definite reference. For physical copresence, as in deixis, people prototypically use demonstratives. For linguistic copresence, as in anaphora, they prototypically use pronouns or definite descriptions. And for community membership, they prototypically use proper names, especially proper nouns. The heuristics also determine in part how people repair inadequate or unsuccessful definite references. The idea is that each repair should strengthen the basis on which mutual knowledge of the referent is established. The copresence heuristics, by spelling out the trade-off between direct evidence and certain auxiliary assumptions, tell how that basis can be strengthened.

And these heuristics require a memory that is organized around diary entries and around communities in which knowledge is universally shared. Currently, the memory assumed in most models of comprehension and production is not organized this way.

What all this suggests is that our views of comprehension and production are in need of reform. We have tried to shatter the illusion that definite reference is simple and self-evident by demonstrating how it requires mutual knowledge, which complicates matters enormously. But virtually every other aspect of meaning and reference also requires mutual knowledge, which also is at the very heart of the notion of linguistic convention and speaker meaning. Mutual knowledge is an issue we cannot avoid. It is likely to complicate matters for some time to come.

NOTES

1 One important caveat here. Often all Ann will be able to check is her belief or assumption or supposition instead of her *knowledge* that t is R. Which propositional attitude is appropriate – knowledge, belief, assumption, supposition, or even some other .term – depends on the evidence Ann possesses and other factors. For simplicity we will use *know* as the general term, but we could replace it with *believe* or certain other terms without affecting our argument.

2 Another way to represent this is as two interreferring statements of this kind:

A and B mutually know that p =_{def.}
> (r)   A knows that p and that r'.
> (r') B knows that p and that r.

In some ways this representation is preferable, for unlike the single self-referential statement, it does not assume that if A knows that A knows that p, then A knows that p. Although this assumption may be justifiable for the verb *know,* it is not so obviously justifiable with *believe* or *assume* or *suppose* in place of *know.*

3 This one-sided mutual knowledge can be represented in a self-referential definition (see note 2) as follows:

A knows that A and B mutually know that $p =_{def.}$
(r) A knows that p and that: B knows that p and that r.

4 "Visibility" is obviously too restricted a term here and should be replaced by "perceptability" to encompass taste, smell, and hearing, as in *where is the awful smell/taste/noise coming from?* This is part of our reason for later using the term "physical" as opposed to "visual" for such cases.

5 It cannot be correct for other reasons either. In Donnellan's (1966) example "Who is the man drinking a martini?", the definite reference *the man drinking a martini* refers to a particular man even if the speaker is mistaken and the man happens to be drinking water; moreover, such a reference will generally succeed (see also Donnellan, 1968). The complication of this sort of misdescription and its relation to mutual belief are thoroughly discussed by Perrault and Cohen (this volume). There is a related problem in deception (see Bruce and Newman, 1978).

6 There may seem to be no real difference between the selective and the progressive checking strategies because for shared knowledge the truth of condition (4), for example, entails the truth of conditions (1), (2), and (3). For shared beliefs or suppositions or other propositional attitudes, however, this entailment no longer holds. If Ann believes that Bob believes that the movie showing at the Roxy tonight is *A Day at the Races,* that doesn't imply that *she* believes it is *A Day at the Races.* In the more general case, these two strategies are distinct.

7 When mutual knowledge is treated as a primitive, it follows that most cases of non–mutual knowledge will require a more complex memory representation than mutual knowledge. As a consequence, they ought to be more difficult to understand. Our Marx brothers scenarios bear out this prediction. Versions 2, 3, 4, and 5 were successively more difficult to understand. The knowledge we had to keep in mind required more and more conditions, and these conditions themselves became more and more complex. The version for mutual knowledge, where Ann and Bob openly discussed the showing of *Monkey Business* at the Roxy, was the easiest to understand. Apparently, the mutual knowledge we had to represent for it was simple.

8 Of course, we must qualify the notion that *everybody* in a community needs to know a thing before it is taken to be mutual knowledge within that community. We can do that informally by replacing *everybody* with *almost everybody,* and *universal* by *almost universal,* or we can do it more formally by introducing parameters that specify the probabilities (see Lewis, 1969, pp. 76–80). This qualification is needed if we want to account for why certain references that are otherwise justifiable on the basis of community membership and community knowledge occasionally fail.

9 Written language, as in books and on signs, we assume, is derivative from spoken language and requires an extended notion of copresence. In *Pride and Prejudice,* for example, Jane Austen assumed her readers would be rational comprehending people who would take in her words serially, as if spoken, etc., etc. She could pretend, in other words, that she was speaking her novel to each reader and that linguistic copresence would be established that way. Signs often rely on an extended notion of physical

copresence as well. For example, *Break this glass to sound alarm* makes sense on a fire alarm, but not pinned to the back of a professor's coat. Nevertheless, we are mindful of the differences between written and spoken language and expect them to complicate the copresence heuristics in various ways.

10 One exception, pointed out to us by a native, is the Highland Scottish use of *himself* as a proper name for the local laird or head of a clan, as in *Himself was angry with Ian today.* Its highly marked form helps to make its proper-name status clear.

11 What an adequate answer to a *who*-question consists of has been taken up by Boer and Lycan (1975) in their paper on "knowing who." They argue that the answer to "Who is X?" is always relative to some purpose and that its ultimate answer is always an attributive use of the definite description. So the *ultimate* answers – and the ultimate repairs – go beyond our paper, which is about referential uses.

REFERENCES

Anderson, J. R. 1976. *Language, Memory, and Thought.* Hillsdale, N.J.: Erlbaum.

Anderson, J. R. 1977. "Memory for Information about Individuals." *Memory and Cognition* 5:430–42.

Anderson, J. R. 1978. "The Processing of Referring Expressions within a Semantic Network." In D. L. Waltz, ed. *Theoretical Issues in Natural Language Processing—2.* New York: Association for Computing Machinery, pp. 51–6.

Anderson, R. C., Pichert, J. W., Goetz, E. T., Schallert, D. L., Stevens, K. V., and Trollip, S. R. 1976. "Instantiation of General Terms." *Journal of Verbal Learning and Verbal Behavior* 15:667–79.

Barenboim, C. 1978. "Development of Recursive and Nonrecursive Thinking about Persons." *Developmental Psychology* 14:419–20.

Boer, S. E., and Lycan, W. G. 1975. "Knowing Who." *Philosophical Studies* 28:299–344.

Bruce, B., and Newman, D. 1978. "Interacting Plans." *Cognitive Science* 2:195–234.

Carpenter, P. A., and Just, M. A. 1977. "Integrative Processes in Comprehension." In D. LaBerge and S. J. Samuels, eds. *Basic Processes in Reading: Perception and Comprehension.* Hillsdale, N.J.: Erlbaum.

Chafe, W. L. 1974. "Language and Consciousness." *Language* 50:111–33.

Christopherson, P. 1939. *The Articles: A Study of Their Theory and Use in English.* Copenhagen: Munksgaard.

Clark, E. V. 1973. "What's in a Word? On the Child's Acquisition of Semantics in His First Language." In T. E. Moore, ed. *Cognitive Development and the Acquisition of Language.* New York: Academic Press, pp. 65–110.

Clark, E. V. 1978. "From Gesture to Word: On the Natural History of Deixis in Language Acquisition." In J. S. Bruner and A. Garton, eds. *Human Growth and Development: Wolfson College lectures 1976.* Oxford: Oxford University Press, pp. 85–120.

Clark, E. V., and Andersen, E. S. 1979. "Spontaneous Repairs: Awareness in the Process of Acquiring Language." Paper presented at the biennial

meeting of the Society for Research in Child Development, San Francisco. March.

Clark, E. V., and Clark, H. H. 1979. "When Nouns Surface as Verbs." *Language,* 55:767–811.

Clark, E. V., and Sengul, C. J. 1978. "Strategies in the Acquisition of Deixis." *Journal of Child Language* 5:457–75.

Clark, H. H. 1977. "Inferences in Comprehension." In D. LaBerge and S. J. Samuels, eds. *Basic Processes in Reading: Perception and Comprehension.* Hillsdale, N.J.: Erlbaum.

Clark, H. H. 1978. "Inferring What Is Meant." In W. J. M. Levelt and G. B. Flores d'Arcais, eds. *Studies in the Perception of Language.* London: Wiley.

Clark, H. H. 1979. "Responding to Indirect Speech Acts." *Cognitive Psychology,* 11:430–477.

Clark, H. H., and Clark, E. V. 1977. *Psychology and Language: An Introduction to Psycholinguistics.* New York: Harcourt Brace Jovanovich.

Clark, H. H., and Haviland, S. E. 1977. "Comprehension and the Given–New Contract." In R. O. Freedle, ed. *Discourse Production and Comprehension.* Norwood, N.J.: Ablex Publishing, pp. 1–40.

Clark, H. H., and Marshall, C. R. 1978. "Reference Diaries." In D. L. Waltz, ed. *Theoretical Issues in Natural Language Processing—2.* New York: Association for Computing Machinery, pp. 57–63.

Clark, H. H., and Sengul, C. J. 1979. "In Search of Referents for Nouns and Pronouns." *Memory and Cognition* 7:35–41.

Cohen, P. R. 1978. "On Knowing What to Say: Planning Speech Acts." Unpublished doctoral dissertation, University of Toronto.

Donnellan, K. S. 1966. "Reference and Definite Descriptions." *Philosophical Review* 75:281–304.

Donnellan, K. R. 1968. "Putting Humpty Dumpty Together Again." *Philosophical Review* 77:203–15.

DuBois, J. W. 1975. "Syntax in Mid-Sentence." In *Berkeley Studies in Syntax and Semantics,* Vol. 1. Berkeley, Calif.: Institute of Human Learning and Department of Linguistics, University of California, pp. III-1–III-25.

Flavell, J. H., Botkin, P. T., Fry, C. L., Wright, J. W., and Jarvis, P. E. 1968. *The Development of Role Taking and Communication Skills in Children.* New York: Wiley.

Grice, H. P. 1957. "Meaning." *Philosophical Review* 66:377–88.

Harman, G. 1977. "Review of *Linguistic Behavior* by Jonathan Bennett." *Language* 53:417–24.

Haviland, S. E., and Clark, H. H. 1974. "What's New? Acquiring New Information as a Process in Comprehension." *Journal of Verbal Learning and Verbal Behavior* 13:512–21.

Hawkins, J. A. 1978. *Definiteness and Indefiniteness: A Study in Reference and Grammaticality Prediction.* London: Croom Helm.

Jespersen, O. 1949. *A Modern English Grammar on Historical Principles, Vol. 7.* Copenhagen: Munksgaard.

Karttunen, L. 1977. "Presupposition and Linguistic Context." In A. Rogers, B. Wall, and J. P. Murphy, eds. *Proceedings of the Texas Conference on Performatives, Presuppositions, and Implicatures.* Arlington, Va.: Center for Applied Linguistics, pp. 149–60.

Karttunen, L., and Peters, S. 1975. "Conventional Implicature in Montague

Grammar." In *Proceedings of the First Annual Meeting of the Berkeley Linguistic Society,* Berkeley, Calif.: Berkeley Linguistics Society, University of California, pp. 266–278.

Kempson, R. M. 1975. *Presupposition and the Delimitation of Semantics.* Cambridge: Cambridge University Press.

Kintsch, W. 1974. *The Representation of Meaning in Memory.* Hillsdale, N.J.: Erlbaum.

Kintsch, W., and van Dijk, T. A. 1978. "Toward a Model of Text Comprehension and Production." *Psychological Review* 85:363–94.

Kripke, S. 1972. "Naming and Necessity." In D. Davidson and G. Harman, eds. *Semantics of Natural Language.* Dordrecht: D. Reidel.

Kripke, S. 1977. "Speaker's Reference and Semantic Reference." *Midwest Studies in Philosophy* II:255–76.

Lasnik, H. 1976. "Remarks on Coreference." *Linguistic Analysis* 2:1–22.

Lesgold, A. M., Roth, S. F., and Curtis, M. E. 1979. "Foregrounding Effects in Discourse Comprehension." *Journal of Verbal Learning and Verbal Behavior* 18:291–308.

Lewis, D. K. 1969. *Convention: A Philosophical Study.* Cambridge, Mass.: Harvard University Press.

Lyons, J. 1975. "Deixis as the Source of Reference." In E. L. Keenan, ed. *Formal Semantics of Natural Language.* Cambridge: Cambridge University Press.

Lyons, J. 1977. *Semantics,* Vols. 1 and 2. Cambridge: Cambridge University Press.

McCawley, J. D. 1979. "Presupposition and Discourse Structure." In C. K. Oh, ed. *Syntax and Semantics,* Vol. 13, *Presupposition.* New York: Academic Press.

Maratsos, M. P. 1976. *The Use of Definite and Indefinite Reference in Young Children: An Experimental Study of Semantic Acquisition.* Cambridge: Cambridge University Press.

Miller, P. H., Kessel, F. S., and Flavell, J. H. 1970. "Thinking about People Thinking about People Thinking about . . . : A study of Social–Cognitive Development." *Child Development* 41:613–23.

Minsky, M. 1975. "A Framework for Representing Knowledge." In P. Winston, ed. *The Psychology of Computer Vision.* New York: McGraw-Hill, pp. 211–77.

Moravcsik, E. A. 1969. "Determination." *Working Papers in Language Universals.* Stanford, Calif.: Stanford University, 1:64–98.

Morgan, J. L. 1978. "Two Types of Convention in Indirect Speech Acts." In P. Cole, ed. *Syntax and Semantics,* Vol. 9, *Pragmatics.* New York: Academic Press, pp. 261–80.

Nunberg, G. D. 1977. "The Pragmatics of Reference." Unpublished doctoral dissertation, City University of New York.

Ortony, A., and Anderson, R. C. 1977. "Definite Descriptions and Semantic Memory." *Cognitive Science* 1:74–83.

Prince, E. F. 1978. "On the Function of Existential Presupposition in Discourse." In *Papers from the Fourteenth Regional Meeting, Chicago Linguistics Society.* Chicago: Chicago Linguistics Society, University of Chicago, pp. 362–76.

Rumelhart, D. E., and Ortony, A. 1977. "The Representation of Knowledge

in Memory." In R. C. Anderson, R. J. Spiro, and W. E. Montague, eds. *Schooling and the Acquisition of Knowledge.* Hillsdale, N.J.: Erlbaum.

Rumelhart, D. E., Lindsay, P. H., and Norman, D. A. 1972. "A Process Model for Long-Term Memory." In E. Tulving and W. Donaldson, eds. *Organization of Memory.* New York: Academic Press, pp. 197–246.

Schank, R., and Abelson, R. 1977. *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures.* Hillsdale, N.J.: Erlbaum.

Schegloff, E. A. 1972. "Notes on a Conversational Practice: Formulating Place." In D. N. Sudnow, ed. *Studies in Social Interaction.* New York: The Free Press, pp. 75–119.

Schegloff, E. A., Jefferson, G., and Sacks, H. 1977. "The Preference for Self-correction in the Organization of Repair in Conversation." *Language* 53:361–82.

Schelling, T. C. 1960. *The Strategy of Conflict.* Oxford: Oxford University Press.

Schiffer, S. R. 1972. *Meaning.* Oxford: Oxford University Press.

Searle, J. R. 1965. "What is a Speech Act?" In M. Black, ed. *Philosophy in America.* London: Allen and Unwin, pp. 221–39.

Searle, J. R. 1969. *Speech Acts: An Essay in the Philosophy of Language.* Cambridge: Cambridge University Press.

Searle, J. R. 1975. "Indirect Speech Acts." In P. Cole and J. L. Morgan, eds. *Syntax and Semantics,* Vol. 3, *Speech Acts.* New York: Seminar Press, pp. 59–82.

Stalnaker, R. C. "Pragmatic Presuppositions." In A. Rogers, B. Wall, and J. P. Murphy, eds. *Proceedings of the Texas Conference on Performatives, Presuppositions, and Implicatures.* Arlington, Va.: Center for Applied Linguistics, pp. 135–48.

Stalnaker, R. C. 1978. "Assertion." In P. Cole, ed. *Syntax and Semantics,* Vol. 9, *Pragmatics.* New York: Academic Press, pp. 315–32.

Sternberg, S. 1966. "High-speed Scanning in Human Memory." *Science* 153:652–4.

Strawson, P. F. 1964. "Intention and Convention in Speech Acts." *Philosophical Review* 73:439–60.

Strawson, P. F. 1974. *Subject and Predicate in Logic and Grammar.* London: Methuen.

Thorne, J. P. 1972. "On the notion 'definite.'" *Foundations of Language* 8:562–8.

Thorne, J. P. 1974. "Notes on 'Notes on "On the notion 'definite.' " ' " *Foundations of Language* 11:111–14.

Townsend, J. T. 1972. "Some Results Concerning the Identifiability of Parallel and Serial Processes." *British Journal of Mathematical and Statistical Psychology* 25:168–99.

Warden, D. A. 1976. "The Influence of Context on Children's Use of Identifying Expressions and References." *British Journal of Psychology* 67:101–12.

Winograd, T. 1972. *Understanding Natural Language.* New York: Academic Press.