

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/246878502>

Mutual knowledge and relevance in theories of comprehension

Article · January 1982

CITATIONS

73

READS

176

2 authors:



Dan Sperber

Central European University

160 PUBLICATIONS 9,982 CITATIONS

SEE PROFILE



Deirdre Wilson

University College London

93 PUBLICATIONS 6,316 CITATIONS

SEE PROFILE

Sperber + Wilson 1982

Mutual Knowledge

edited by

N. V. Smith

*Department of Phonetics and Linguistics
University College, London*

1982



ACADEMIC PRESS

A Subsidiary of Harcourt Brace Jovanovich, Publishers

London · New York
Paris · San Diego · San Francisco
São Paulo · Sydney · Tokyo · Toronto

Chapter 2

Mutual Knowledge and Relevance in Theories of Comprehension¹

Dan Sperber and Deirdre Wilson

I. INTRODUCTION

The main aim of pragmatic theory is to explain how successful communication is possible, and in particular, how utterances are understood. Understanding an utterance involves recovering the proposition it expresses and drawing certain inferences based on this proposition as premise. The difficulty lies in explaining not how some arbitrary proposition is discovered and some random inferences drawn, but how the intended content and intended implications are recovered: that is, how comprehension is achieved.

Comprehension is a function of the context: that much is uncontroversial. But what does the context consist of? How does the hearer exploit it in discovering the intended content and implications of an utterance? Various answers to these questions have been proposed, but all of them are tentative at best.

It is compatible with much of the current literature to envisage three distinct mechanisms: one to determine the context involved in the comprehension of an utterance, a second to determine the content on the basis of the context and of the linguistic properties of the utterance, and a third to draw the intended inferences on the basis of the content and the context.

As regards the context, some recent work suggests that it is restricted to the mutual knowledge, beliefs and suppositions of speaker and hearer, where mutual knowledge is knowledge that is not only shared,

but known to be shared, and known to be known to be shared, and so on. On this approach, the identification of mutual knowledge is a major factor in every aspect of comprehension, and one of the most urgent goals of pragmatic theory is to explain how it is achieved.

We would like to develop three main arguments against this approach. First, the identification of mutual knowledge presents problems which, contrary to the predictions of the mutual knowledge framework, do not give rise to corresponding problems of comprehension. Secondly, mutual knowledge is not a sufficient condition for belonging to the context: a proposition may be mutually known without being part of the context. Thirdly, it is not a necessary condition either: a proposition may belong to the context without being mutually known.

We shall then suggest an alternative approach. We shall argue that there is a single principle which simultaneously determines context, content and intended inferences, with no appeal to mutual knowledge.² The fact that some knowledge is considered mutual is generally a result of comprehension rather than a precondition for it. Hence mistakes in comprehension are much more likely to cause a wrong assessment of mutual knowledge than the other way around.

II. SOME QUESTIONS FOR THE MUTUAL KNOWLEDGE FRAMEWORK

A. Do Problems in Identifying Mutual Knowledge Cause Problems in Comprehension?

In order to understand an utterance, the hearer has to bring to bear certain items of background information not specifically mentioned in the utterance. For instance, in order to understand the utterance "I didn't", the hearer must be able to identify some activity, not specified in the utterance, which he can take the speaker to be saying he did not engage in. The context, as generally understood, is the background information that can be brought to bear on comprehension.

What is the extent of the context intrinsically involved in the comprehension process? (For the contrast between "intrinsic" and "incidental" context, see Clark and Carlson, forthcoming.) The weakest hypothesis would be that all the information the hearer possesses can be brought to bear on comprehension: that is, that the context is co-extensive with the hearer's memory. However, a more restrictive hypothesis

is generally favoured. It is argued, for instance by Clark and Carlson (forthcoming) that the intrinsic context can be straightforwardly identified with the mutual knowledge, beliefs and assumptions of speaker and hearer: that is, with what has been referred to elsewhere in the literature as their common ground (Stalnaker 1974; Karttunen and Peters 1975). There are formal arguments for this identification of context with mutual knowledge or common ground. We shall discuss them later. But to begin with, there are problems with the identification of mutual knowledge itself.

The concept of mutual knowledge was first introduced as part of the philosophical analysis of speaker-meaning, utterance-meaning, convention and other meaning-related concepts (Schiffer 1972, 30-42; Lewis 1969, 52-60). For those interested in constructing an empirical pragmatic theory, the question is not whether these analyses are philosophically adequate, but whether they have any psychological correlates. Here an immediate problem arises. If mutual knowledge is to play a role in the real-time production and comprehension of utterances, it must be very easily identifiable: there must be some straightforward method by which a speaker and hearer who both know a given proposition can discover that they mutually know it. But at first sight, it is hard to see how such a method could exist.

Mutual knowledge is knowledge of an infinite set of propositions. By the usual definitions, a speaker S and an addressee A mutually know a proposition P if and only if:

- (1) S knows that P.
- (2) A knows that P.
- (3) S knows (2).
- (4) A knows (1).
- (5) S knows (4).
- (6) A knows (3).

... and so on *ad infinitum*.

How does A discover that he has the requisite mutual knowledge for understanding an utterance? Assuming that he cannot compute an infinite set of propositions in a finite amount of time, and that the possession of mutual knowledge is not self-evident (since one can be mistaken about it), the problem is to find a finite procedure for distinguishing mutual knowledge from knowledge that is not mutual.

Clark and Marshall (1981) argue that mutual knowledge can be identified by a finite inductive procedure. They argue that from the

psychological point of view, mutual knowledge must be a simple, unanalysable concept, whose logical consequences do not have to be computed in order to establish its applicability. Instead, its applicability is established by invoking more or less adequate inductive evidence (see also Lewis 1969, 52–58; Schiffer 1972, 33–36). Clark and Marshall classify this evidence according to its possible sources: physical co-presence, linguistic co-presence and community membership.

Physical co-presence provides the strongest evidence for mutual knowledge. It involves the presence of two people, S and A, at an event which gives them direct empirical evidence for a certain proposition, and for the fact that both of them have this evidence. For example, if S and A are facing each other across a table with a bowl of fruit between them, then with certain minimal assumptions about rationality and powers of observation, both would be justified in concluding that they had mutual knowledge of (7):

(7) There is a bowl of fruit on the table between S and A.

This piece of mutual knowledge could be stored in memory as such, or reconstructed at a later date given mutual knowledge of its retrievability.

Linguistic co-presence involves the co-presence of S, A and an utterance which expresses or implies a certain proposition. For example, if S and A are standing together waiting for a train to Oxford when the station-announcer makes the announcement in (8), they could reasonably infer mutual knowledge of (8):

(8) The 6:00 train for Oxford will leave from platform 6.

Again, this could be stored at the time as a piece of mutual knowledge, or reconstructed later if it is mutually known to be retrievable.

Finally, if S and A establish that they belong to the same community or group, they can reasonably assume mutual knowledge of all propositions normally known by group members. For example, if it is established that they have both been exposed to a recent Heineken's advertizing campaign, they can assume mutual knowledge of (9):

(9) It has been claimed that Heineken's refreshes the parts other beers cannot reach.

The above remarks about storage and retrievability still apply.

Much of the literature on mutual knowledge explicitly assumes that all knowledge is evidenced. Clark and Marshall implicitly assume that all evidence for mutual knowledge is ultimately physical. In their

terms, linguistic co-presence is simply physical co-presence at an acoustic (or visual) event, and community co-membership has to be established through events of physical or linguistic co-presence. Clearly, one might need quite lengthy chains of evidence to connect a particular item of mutual knowledge to the physical evidence which supports it, and in a largely inductive framework, each step in the chain may go well beyond the data. Moreover, even in the case of physical co-presence, which provides the strongest evidence for mutual knowledge, powerful systems of auxiliary assumptions may have to be used in establishing a conclusion. For example, although the presence of a visible object may be self-evident to any observant, sighted person, a considerable number of auxiliary assumptions would be needed to establish mutual knowledge of it under a particular description: as a cricket match, or the coronation of Queen Elizabeth II, or even a bowl of fruit.

Given these difficulties in the identification of mutual knowledge, if identification is a prerequisite to comprehension, an obvious empirical prediction follows. The problems that arise in assessing mutual knowledge should cause corresponding problems in comprehension. In particular, when the evidence required for mutual knowledge goes well beyond straight physical co-presence, there should always be some room for doubt in the hearer's mind about whether he has correctly understood. This is not borne out by introspective evidence. It seems much easier to understand an utterance than it does to assess mutual knowledge. Moreover, an utterance such as (10) does not seem harder to understand than (11), although the mutual knowledge involved is much harder to establish inductively:

(10) I dislike the eldest Brontë sister.

(11) I dislike that girl over there.

Of course, such unsystematic introspective evidence is not enough to settle the issue. But it should at least cast doubt on the identification of context with common ground, and make it desirable to look for an alternative framework.

B. Is Mutual Knowledge a Sufficient Condition for Belonging to the Context?

As we have said, using the context in comprehension involves retrieving specific items of information. For instance, if the speaker says:

(12) I am a moslem. I don't drink alcohol,
the hearer's task is to retrieve the background information in (13):

(13) Moslems are forbidden to drink alcohol.

Given (13), (12) can be understood as both stating a fact and explaining it. Without (13), (12) would seem to state two unrelated facts, and the speaker's intention would not be understood.

In order to retrieve propositions such as (13), that may have a bearing on comprehension, some search of the context has to take place. However, given the speed of actual comprehension and the size of the common ground shared by members of the same community, it is not really possible that every proposition in the common ground is parsed and checked for a possible role in the comprehension of every utterance. For instance, in the case of (12), although speaker and hearer may have much more mutual knowledge of Christianity than of Islam, it seems likely that the hearer will search only his knowledge about Islam. Some other criterion than just belonging to common ground must be used to determine the context actually searched, reducing it to a manageable size.

Apart from manageability, there is another reason for taking a much more restrictive view of the context actually searched. The role of context in comprehension that has probably been the most discussed is in the assignment of reference. In interpreting an utterance which contains the referring expression *the door*, the hearer has to search the context in order to identify the actual door being referred to. In many cases, speaker and hearer mutually know of the existence of more than one door. However, not all of these will be actively considered as possible referents. This again suggests that the actual context used in comprehension is much smaller than the common ground.

When Clark and Marshall and others equate the context with mutual knowledge or common ground, what they have in mind is not the context actually searched but a potential context of which the actual one is a small subset. That is to say, they might grant that belonging to mutual knowledge is not a sufficient condition for a proposition's being part of the actual context, but claim that it is a necessary condition and, furthermore, the only necessary condition. If this claim were correct, it would make sense to talk of mutual knowledge as providing a potential context, smaller than the subject's encyclopaedic memory and larger than the actual context. There is no evidence or argument in the literature to show that belonging to mutual knowledge is the only necessary

condition, so that the case for identifying potential context with mutual knowledge is anyway incomplete. However, there are arguments to show that it is a necessary condition, and to these we now turn.

C. Is Mutual Knowledge a Necessary Condition for Belonging to the Context?

Given that it is, if not downright impossible, at least fairly cumbersome to establish mutual knowledge, and given that mutual knowledge is not sufficient for determining the context in which an utterance has to be understood, why bother to establish it at all? Clark and Marshall (1981), while acknowledging that "it is likely to complicate matters for some time to come", argue that "mutual knowledge is an issue we cannot avoid", because

virtually every . . . aspect of meaning and reference . . . requires mutual knowledge, which also is at the very heart of the notion of linguistic convention and speaker meaning.

The argument for this admittedly expensive claim consists in showing that, as long as background knowledge is only shared to some degree but is not fully mutual, comprehension is not guaranteed. The argument bears a strong structural similarity to earlier philosophical arguments that mutual knowledge plays a necessary role in the analysis of speaker-meaning (see, for example, Schiffer 1972, pp. 30-42).

Clark and Marshall's version of this argument involves definite reference, and revolves around episodes such as the following:

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?" (p. 5)

Here, Bob is likely to misunderstand Ann's reference to the movie showing at the Roxy, because although Ann knows that Bob knows that *Monkey Business* is the film showing, she doesn't know whether Bob knows that she knows it too. More elaborate episodes involve further

degrees of sharedness, and since there is no intrinsic limit on this process of elaboration, it follows that full-scale *mutual* knowledge is necessary in order to be sure that reference will be properly understood.

We believe that the unnaturalness of these examples is not accidental, and that it should have warned psychologists against following the lead of the philosophers in this area. In real life, if any such unnaturally complex situation arose, either the hearer would ask for clarification, or as likely as not, misunderstanding *would* occur.

The very strength of the formal argument should cast doubt on its empirical relevance. Its proponents see language as governed by a set of conventions, grammars as constructed out of sets of conventions, and conventions themselves as analysable in terms of mutual knowledge (Schiffer 1972, pp. 118–160; Lewis 1969, pp. 60–68; Clark and Carlson forthcoming, pp. 22–23). One issue here is whether grammars are best seen as collections of conventions. We see no reason for thinking that they are, but we will not dispute it here (see Smith and Wilson, 1979, pp. 14–21). But what is true is that the mutual knowledge argument formally applies as much when it comes to determining what language has been used as it does to determining what context was intended.

If someone addresses me with an utterance that has all the appearances of being an utterance in my own native language, how can I know for sure that it is one? It could belong to some entirely different language and, by coincidence, sound exactly like an utterance of mine.³ Less fantastically, it could belong to some language that has the same phonological structures but not the same meanings as mine: it could be some distant dialect, for example. To be sure that this is not happening, I have to know that the speaker knows my native language, and knows that I know it, and knows that I know that he knows it, and so on. Moreover, it is not just mutual knowledge of my native language that would be necessary, but also of the fact that this particular utterance was made in it.

On the other hand, it is patently obvious that much, if not all, of verbal communication takes place without these conditions being satisfied. For instance, an English speaker visiting a foreign country, who walks up to some passer-by and asks “Do you speak English?” would not do so if they had mutual knowledge of English, and cannot safely do so unless they have. In other words, the tourist takes a risk. In fact we all take risks, whenever we engage in verbal communication. At

this moment, we are taking the risk of being misunderstood, you are taking the risk of misunderstanding us, and yet we proceed.

What this suggests is that the formal argument is irrelevant to actual comprehension. It leaves out a simple fact: we don't need to be *sure* that a remark is, say, in English, but only to have sufficient ground for assuming that it is. The fact that it could be an utterance in English is, in almost every case, sufficient reason for thinking that it is one. It is not just that we do not need to be sure: in fact, we *could not* be sure, since mutual knowledge itself cannot be established with absolute certainty.

Similarly, it is probably correct that we could not be sure of successful disambiguation, reference assignment, recovery of the intended inferences, and so on, without mutual knowledge of the context. But what this establishes is the trivial point that we cannot be sure, not the controversial point that mutual knowledge of the context is a necessary part of comprehension. It could still be, as Clark and associates might want to argue, that the strategy of comprehension consists in aiming at certainty and trying to get as close as possible to it. This would imply that the subject takes all feasible steps which would be necessary for achieving certainty, although they could not be sufficient. This, however, is an implausible strategy: it takes only cognitive benefits into account and ignores processing costs.

Going through all the problems involved in assessing mutual knowledge in order to be sure of understanding is like paying a heavy premium to an insurance company which cannot be trusted to cover the risk insured. It is generally not worth it. The only cases where a genuine effort is made to establish mutual knowledge of the meaning, reference and implications of texts are legal documents and treatises, where the risk involved in misunderstanding is so great that the cost of reducing it is acceptable. Even then, as lawyers well know, mutual knowledge is by no means always achieved.

The formal argument that mutual knowledge is a necessary condition for comprehension applies only to perfect comprehension, and not to the imperfect form which is felt to be quite sufficient in daily life. Once this is realized, it is easy to see the many counter-examples to the mutual knowledge approach. For instance:

Ann believes that Bob does not know which film is showing tonight at the Plaza. But to annoy him, she nevertheless asks: “Have you ever seen the film showing at the Plaza tonight?”. It so happens that Bob knows that the film in

question is *Wuthering Heights* and knows that Ann believes that he does not know it. He answers nevertheless "Yes, I have". She understands correctly that he has seen *Wuthering Heights*, and infers that she was wrong to believe that he did not know which film was showing at the Plaza. The lack of mutual knowledge leaves room for doubt: she cannot be sure that Bob is referring to *Wuthering Heights*, but she is led to think (correctly) that he is. As a result, but not a condition, of this act of comprehension, it is mutually assumed to be known what film is showing tonight.

Or:

Bill, while travelling in Southern Europe, offers a cigarette to a peasant whom he believes to be ignorant. The peasant answers "No thank you, I have read the latest statistics". Bill is surprised, but understands correctly that the peasant wants him to take as part of the context the fact that the latest statistics show that smoking is hazardous to one's health, and to infer from that context and the peasant's answer the reason why his offer of a cigarette is declined. (Of course, Bill cannot be sure that this is what the peasant meant.) As a result of this act of comprehension, the fact that smoking is hazardous becomes mutually assumed to be known.

As these two examples show, a proposition can be included in the context although it is not part of the common ground, and indeed although it is believed not to be shared at all. It can be used in establishing reference or intended inferences. As a result of its having been so used, and only then, it will be assumed to be mutually known. Hence, a model of comprehension need not have a device for establishing the common ground as one of its sub-parts: on the contrary, the model as a whole should constitute one of the elements of a theory of how the common ground is established.

III. AN ALTERNATIVE FRAMEWORK

A. Relevance

The following objection could be made to the arguments of the last section. In identifying the language in which a remark is made, mutual knowledge is, in practice, unnecessary for a simple reason. If an utterance sounds like a sentence of English, it is safe enough to assume that it is one, because there are in almost every case no real alternatives. So although mutual knowledge would be necessary for an absolute iden-

tification, a reasonable identification can be achieved without it. In understanding a given utterance, on the other hand, there *are* alternatives. Almost every sentence is semantically ambiguous; it contains referring expressions which range over a large, perhaps infinite, domain; from the logical point of view, it has an infinite set of contextual implications in each of an indefinite number of possible contexts. Even if infinities can be trimmed down to finite numbers, there is always, for any given utterance, a large range of interpretations compatible with its semantic, referential and inferential properties. Because of this, there has to be some mechanism, whether simple or complex, which singles out the intended interpretation. What little understanding we have of this mechanism at present seems to rest on the exploitation of knowledge that is not just shared but mutual. Unless some alternative framework is provided, criticism of the mutual knowledge framework might force one to amend it, but surely not to abandon it.

We believe that an alternative framework can be developed. The basic insight on which it rests was suggested by Grice (1967; see Grice 1975, 1978 for published parts). Comprehension, he argued, involves not only a particular utterance and a particular context, but also the presumption that the speaker has tried to conform to some general standards of verbal communication. This presumption is used as a guide to the intended interpretation. The question is whether this guide is good enough for comprehension to proceed without a prior assessment of mutual knowledge: this in turn depends on how constraining the standards governing verbal communication turn out to be.

The standards actually proposed by Grice (his now famous Cooperative Principle and maxims of conversation) are not constraining enough for these purposes. They fail for two reasons: first, they are far too vague: it has never been fully specified exactly what their content is, nor exactly how they are supposed to function. However, this defect could, at least in principle, be overcome. Secondly, and more seriously, Grice's model relaxes some crucial constraints on the comprehension process by allowing the literal meaning of figurative utterances to act as no more than a loose set of hints at the intended message, rather than being a necessary part of the message itself. As a result, the number of available interpretations for every utterance is increased in an uncontrollable way (see Sperber and Wilson, 1981; Wilson and Sperber, 1981, for further discussion of these points).

We would like to propose a constraint that will be much more restric-

tive than Grice's. We shall argue that speakers try, and hearers expect them to try, to meet a single general standard in producing an utterance: a standard of relevance. We have developed the argument more technically and in much greater detail elsewhere (see Sperber and Wilson, forthcoming). Here we shall give only a very informal presentation; perhaps not enough to demonstrate its validity, but enough, we hope, to show that it is possible to conceive of an alternative to the mutual knowledge framework.

Thinking in general, and verbal comprehension in particular, involves drawing inferences. If inferring in these cases consisted in applying the rules of standard logic to some set of premises, an infinite process would take place: an infinite set of inferences would be drawn, most of them of no psychological interest whatsoever. For instance, from two premises 'P' and 'Q', conclusions such as 'P and P', 'P and Q', 'Q and Q', 'P or P', 'P or Q', 'Q or Q', 'P and P and P and P and P or P and Q or Q or Q' can be derived. Clearly, when not doing exercises in formal logic, people do not waste their time deriving these trivial implications. Instead, they concentrate on non-trivial deductions, of which the following are examples:

(14) Slivovitz is an alcohol
 Omar is a moslem
 Moslems don't drink alcohol

 Omar does not drink slivovitz.

(15) The bank closes at five
 It is half past five

 The bank has closed.

An adequate model of inferential processes must likewise distinguish between the infinite set of trivial inferences and the finite set of non-trivial inferences which can be drawn from a finite set of premises, and draw only the latter. In Sperber and Wilson (forthcoming) we provide a characterization of such a model. Here we shall assume the problem solved. We assume, then, that it is possible to compare two sets of premises for the number of non-trivial implications they have. We shall maintain that this is a crucial factor in assessing the relevance of an utterance.

A new utterance in a given context makes it possible to draw infer-

ences which would not be available if either the context or the utterance were missing. We shall call these *contextual implications*. A contextual implication of an utterance is a non-trivial logical implication derivable not from the content of the utterance alone, nor from the context alone, but only from context and content combined. As an illustration, consider the following (attested) exchange:

- (16) a. Flag-seller: Would you like to buy a flag for the Royal National Lifeboat Institution?
 b. Passer-by: No thanks, I always spend my holidays with my sister in Birmingham.

Not everyone finds the response in (16b) immediately comprehensible. In order to understand it fully, the hearer has to supply (at least) the premises in (17), and derive the conclusion in (18):

- (17) a. Birmingham is inland.
 b. The Royal National Lifeboat Institution is a charity.
 c. Buying a flag is one way of subscribing to a charity.
 d. Someone who spends his holidays inland has no need of the services of the Royal National Lifeboat Institution.
 e. Someone who has no need of the services of a charity cannot be expected to subscribe to that charity.

(18) The speaker of (16b) cannot be expected to subscribe to the Royal National Lifeboat Institution.

In our terms, (18) is a contextual implication of (16b) in a context which contains (17). It follows from (16b) and (17) taken together, but from neither (16) nor (17) in isolation from each other.

What is interesting about (16b) from our point of view is the intuitive connection it reveals between being able to derive the contextual implications of an utterance and being able to see its relevance. Those who fail to see the relevance of (16b) at first sight are precisely those who have failed to derive the contextual implication in (18), and anyone who sees this implication will concede the relevance of (16b). This suggests that having contextual implications in a given context is a necessary and sufficient condition for relevance, and can be used as the basis of a definition of relevance.

However, although we could use this definition to formulate a more explicit version of Grice's maxim of relevance, there would be little point in doing so. We are looking for a constraint powerful enough to select at most a single interpretation for any utterance in context. Clearly, the mere requirement that an utterance should have some con-

textual implications, have some degree of relevance, is not powerful enough to do this. We would like instead to investigate the possibility that relevance is not a simple binary concept, but a matter of degree; that one can assign degrees of relevance to possible interpretations, so that speakers and hearers might be conceived of as operating not by a standard of simple relevance, but by a standard of maximal relevance.

B. Degrees of Relevance

The idea of maximal relevance might be usefully approached by an analogy. Consider the measurement of productivity. A firm with output of any value, however small, will be productive to some degree. However, when it comes to measuring productivity, it is not the value of output alone that must be taken into account, but the ratio of output to the value of capital and labour input used in producing it. Of two firms which produce the same output, it will be the one with the smaller input that is the most productive; and of two firms with the same input, it will be the one with the greatest output that is the most productive.

Similar remarks apply to the measurement of relevance. An utterance with any contextual implications, however few, will be relevant to some degree. However, when it comes to comparing degrees of relevance, of different utterances in the same context or the same utterance in different contexts, the number of contextual implications derivable is not the sole factor to be taken into account. Degrees of relevance depend on a ratio of input to output, where output is number of contextual implications, and input is amount of processing needed to derive these contextual implications; by "amount of processing" we mean some function of time and degree of attention expended. Of two utterances that take the same amount of processing, it is the one with most contextual implications that will be the more relevant; and of two utterances which have the same number of contextual implications, it is the one which takes the least amount of processing that will be the more relevant.

To illustrate, compare utterances (19)–(21) in a context consisting of (22a–c):

- (19) Susan, who has thalassemia, is getting married to Bill.
 (20) Susan is getting married to Bill, who has thalassemia.
 (21) Susan, who has thalassemia, is getting married to Bill, and 1967 was a very good year for Bordeaux wine.

- (22) a. People who are getting married should consult a doctor about possible hereditary risks to their children.
 b. Two people both of whom have thalassemia should be warned against having children.
 c. Susan has thalassemia.

In this context, both (19) and (20) carry the contextual implication that Susan and Bill should consult a doctor, but (20) also carries the implication that Susan and Bill should be warned against having children. The sentences in (19) and (20) are almost identical in linguistic and lexical structure. Suppose that processing involves identifying the proposition expressed by the utterance, computing its non-trivial implications, and matching each of these against the propositions in the context to see if further non-trivial implications can be derived. Then (19) and (20) should take roughly equal amounts of processing. In this context, since (20) yields more contextual implications than (19), with the same amount of processing, it should be more relevant than (19), and this seems intuitively correct. By contrast, (19) and (21) have the single contextual implication that Susan and Bill should consult a doctor. (21) is linguistically more complex than (19). On the above assumptions about processing, (21) will thus require more processing and be predicted as less relevant in the context; again, this prediction seems to be intuitively correct.

Given this characterization of relevance in terms of number of contextual implications and amount of processing involved in deriving them, we can spell out what, we suggest, is the single principle governing every aspect of comprehension, the *principle of relevance*:

The speaker tries to express the proposition which is the most relevant one possible to the hearer.

In ordinary circumstances, the hearer assumes that the speaker has not only tried to be as relevant as possible, but has also succeeded. The hearer therefore selects, from all the propositions (i.e. combinations of sense and reference) that the utterance could express, the most relevant one, and assumes that it is the one intended by the speaker.

However, the claim that the principle of relevance governs comprehension in general and disambiguation in particular has little empirical import as long as the context used in comprehension is not specified, and it is to this that we now turn.

C. Choosing a Context: Simplified Version

Most pragmatic accounts assume that the context for the comprehension of a given utterance is fixed in advance, and undergoes no more than minor adjustments during the comprehension process: for example, by the addition of Gricean conversational implicatures (McCawley, 1979 is an interesting exception). We want to argue, on the contrary, that the search for the interpretation on which an utterance will be most relevant involves a search for the context which will make this interpretation possible. In other words, determination of the context is not a prerequisite to the comprehension process, but a part of it. It proceeds, we suggest, as follows.

There is, to begin with, an initial context which consists of the interpretation of the immediately preceding utterance in the conversation or in the text. The hearer attempts an interpretation in this context by looking at what contextual implications can be derived in it. If these are lacking or not considered sufficient to satisfy the principle of relevance, the context can be expanded several times, in three different directions. The hearer can add to the context what he remembers of utterances further back in the conversation (or in previous exchanges with the same speaker). He can add encyclopaedic knowledge which is attached in his memory to the concepts present in the utterance or in the context: for instance, in the example (19)–(21) knowledge of thalassaemia, of Susan and of marriage is present in the context (22). Or he can add to the context information about whatever he is attending to at the same time as the conversation is taking place: for example, information about a football match that speaker and hearer are watching together. The hearer does not have to worry at this stage whether the additions he is making to the context belong to the common ground or not.

Each expansion of the context creates new possibilities of deriving contextual implications. On the other hand, these extensions involve an ever-increasing cost in amount of processing and, in this respect, *diminish* relevance. As a result, if an utterance is not sufficiently relevant in the initial context or a minimally extended one, it is unlikely that its relevance will be increased by further extensions of the context, even though more contextual implications may be found.

The expectation that greater relevance can be achieved by expanding the context, and with it the hearer's readiness to process further, vary with the type of discourse. In ordinary conversation, the time

spent on a given utterance is rarely more than the duration of the utterance itself, and the degree of attention remains relatively low. On the other hand, a believer reading a sacred text tends to take for granted that more time spent, greater attention given, will always lead to an increase in relevance.

Generally speaking, we would suggest that the amount of processing tends to remain roughly constant throughout a stretch of discourse. If this is true, one should expect cases of over- or under-processing to occur, and indeed they do. For instance, readers and writers of scholarly works will know that the point of a dense paragraph, suitable for a professional paper but appearing without warning in the middle of textbook prose, might be missed because of underprocessing, even by readers who would have no problem understanding it otherwise. Conversely, a paragraph in textbook style appearing in a technical paper will get overprocessed and be felt to be insufficiently relevant although, if the whole paper had been written in that style, no problem of relevance need have arisen.

Given the cost of expanding the context and the lack of flexibility in amount of processing, the search for an adequate context tends to remain within a predictable and generally narrow domain. In trying to maximize relevance, the speaker must adapt to this fact, and the hearer can assume that he has. Hence, as a direct result of the principle of relevance, the context is kept down to a manageable size. Restricting it to material from the common ground is unnecessary in this respect at least.

But of course the main argument for restricting context to the common ground has to do with ensuring comprehension. We shall now show that, in this respect too, the principle of relevance makes this restriction unnecessary.

In order to feel confident that his utterance will be adequately relevant to the hearer, the speaker must have grounds for thinking that the hearer has an easily accessible context in which a sufficient number of contextual implications can be derived. One good reason for believing this might be that the required context is part of the common ground; for instance it consists of the interpretation of the three immediately preceding utterances. But this is by no means the only possibility.

The speaker may have grounds for believing that the hearer has access to the required context, without even knowing what this context consists of. For instance, if someone walks up to you in the street and

asks "What time is it?", you assume that the answer you give is relevant to him, that is, that it has a number of contextual implications, without knowing at all what they are and what the context may be.

To take another example, the speaker may know some football results, say, which the hearer does not yet know of, and may believe that the hearer is generally interested in football results. The speaker can then assume that the hearer has a rich enough, easily accessible background of information in the context of which the information being provided will yield many contextual implications. What exactly the context consists of, what these implications are, the speaker does not need to know in order to act in accordance with the principle of relevance.

In neither of these examples were the contextual implications drawn by the hearer specifically intended by the speaker. They were drawn not as part of comprehension proper, but as part of a broader type of processing. However, what our theory of relevance implies is that one of the speaker's intentions (and a crucial one) is that the hearer, by recognizing the speaker's intentions, should be made capable of going beyond them and of establishing the relevance of the utterance for himself. This general intention of being relevant gives the crucial guide to recovery of the meaning, references and inferences (if any) specifically intended by the speaker. A successful act of comprehension (which is what is aimed at by both speaker and hearer) is one which allows the hearer to go beyond comprehension proper.

The speaker thus intends the hearer to draw, or at least to be able to draw, a number of inferences from his utterance. But none of these inferences need be specifically intended. Those inferences which use as premises only the utterance and propositions from the common ground may be presumed to become common ground too, and to be so used in future exchanges. However, that makes them "authorized" rather than "intended" inferences. For example, given the utterance in (23), if (24) is common ground, then presumably the inference in (25) also becomes common ground; this is so even though in uttering (23), the speaker need not have specifically intended that the hearer should infer (25):

(23) Bob is in love with Ann.

(24) Ann is a nuclear physicist.

(25) Bob is in love with a nuclear physicist.

There are two kinds of case where the speaker must assume that a *specific* piece of information will be included in the hearer's context:

cases of definite reference on the one hand, and of intended inference on the other.

In cases of definite reference, the hearer must find in the context one referent for each referring expression. In order to do so, he may have to expand the context (McCawley 1979 develops a similar suggestion). Often enough, the linguistic form of the referring expression gives a clue to the direction in which the extension is to take place: anaphora suggests going back in discourse, proper names suggest a look in the encyclopaedia, and deictics suggest a look around (compare Clark and Marshall's parallel remark, 1981, 41-44). The hearer may then test the assumption that the referents found are the ones intended. If the resulting proposition turns out to be relevant as expected, the assumption will be upheld. Otherwise the context will have to be expanded so as to include other possible referents. When, either from the outset or as a result of expansions, the context contains several possible referents for one referring expression, the principle of relevance will determine the intended one. For example, if pointing at a group of five boys, one of whom is crying, the speaker says "He has just been scolded", the one crying will be selected as the referent of "he". More contextual implications can be derived about this boy than about the other boys by including in the context the fact that he is crying and general background knowledge about crying and scolding.

In cases of intended inference, relevance depends crucially on some specific contextual implication without which the other implications (or at least many of them) cannot be derived. Then, given the principle of relevance, the hearer must assume that the speaker has specifically intended him to draw that inference. For example, consider the following dialogue:

(26) Ann: Will you have a glass of brandy?

Omar: You know I am a good moslem.

If Ann knows that brandy is alcohol and that good moslems do not drink alcohol, she can infer (27):

(27) Omar will not have a glass of brandy.

She can also infer that Omar intended her to draw that specific inference, without which his utterance would not be relevant.

It may happen both in the case of reference and in the case of intended inference, that the hearer can quite easily see what contextual premise would be needed to identify the referent or to draw the intended inference; however, this premise may happen not to be included

in or entailed by what he knows. In such cases, the hearer is entitled to infer from the fact that the utterance has been made, together with the principle of relevance, that the speaker took for granted that the hearer would include this premise in the context for comprehension.

Consider, for example, the utterance in (28) and the contextual premise in (29), which is needed to establish the relevance of (28):

(28) I have read John's novel. The character of Eliza is so moving!

(29) There is a character called Eliza in John's novel.

If the hearer of (28), not having read John's novel, does not know (29), he is entitled to infer it on the basis of his presumption that the speaker of (28) has spoken relevantly. This provides the hearer with a referent for the proper name "Eliza".

Or consider the dialogue in (30) and the contextual premise (31):

(30) Ann: Did you like the book you were reading?

Bob: I don't much like science fiction.

(31) The book Bob was reading is a book of science fiction.

If Ann did not know (31), she can infer it from the fact that without it, Bob's answer would be irrelevant.

The fact that missing premises will be supplied in this way by the hearer can be exploited by the speaker with rhetorical intent. Suppose, for instance, that it was mutual knowledge between Bob and Ann that the book Bob was reading was Chomsky's *Syntactic Structures*. Then Bob's reply to Ann would suggest to her that Bob took for granted that she already believed, or would have no difficulty in accepting, that *Syntactic Structures* is science fiction.

With or without rhetorical intent, these missing premises correspond to Grice's conversational implicatures, and can be derived from the principle of relevance alone, without recourse to the other Gricean maxims.

D. Choosing a Context: Some Refinements

Up to now, we have been considering a simplified version of a model of comprehension based on a single principle of relevance. In this simplified version, the hearer assumes not only that the speaker has tried to make his utterance maximally relevant from the hearer's point of view, but that he has succeeded. In order to do this, the speaker must make some more or less specific assumptions about the contextual propositions that the hearer has access to or can infer. In the simplified

model, the hearer assumes that the speaker's assumptions were correct. It seems to us that most of the time in real life, this would be enough to ensure successful comprehension.

People are fairly accurate in the assumptions they make about what others know. Also, in the case of speech, the speaker can make sure that his assumptions are correct by directly expressing any propositions that he is not sure the hearer would be able to add to the context by himself. Moreover, if the speaker has been significantly wrong in his assumptions, what is likely to happen is not that the hearer will understand something other than the intended propositions: it is rather that the hearer will fail to arrive at a plausible interpretation at all, and will, if he cares enough, ask for repair. It takes quite a coincidence for an utterance to have among its possible interpretations one which the speaker thought would be the most relevant, and a genuinely different one which *is* the most relevant for the hearer. As we noted above, the fact that some utterance closely corresponds to a sentence of English is strong evidence that it was intended as one. Similarly, though to a lesser extent, the fact that one interpretation of an utterance stands out by its greater relevance is strong evidence that it is the intended one.

Nevertheless, for an adequate model of comprehension, we need the principle of relevance as stated: the speaker *tries* to express the most relevant proposition; rather than a simplified version: the speaker expresses the most relevant proposition. Otherwise the model would predict a number of mistakes in comprehension that do not occur in any systematic way.

Take, first, the simple case where the proposition intended by the speaker is completely irrelevant because the hearer already holds this proposition to be true. For instance, the hearer already knew the football results that the speaker was trying to inform him of. In the simplified model, the hearer should fail to identify the intended proposition, and would probably believe that he has not understood. In fact, the full principle of relevance provides the criterion for comprehension: the intended interpretation is generally the only one that the speaker might have thought would be maximally relevant to the hearer, and hence is the only one compatible with the principle of relevance.

A more interesting case is when there is some proposition easily accessible to the hearer which could enormously increase the relevance of the utterance for him but which, being unknown to the speaker, could not possibly have been intended to play a role in the interpretation. In

the simplified model, the hearer should choose the interpretation which maximizes relevance for him, and should thus misunderstand the speaker's intentions and the utterance itself.

However, the full principle of relevance provides a natural check on the conclusion that a certain background proposition was intended to be used. Suppose that to achieve a particular interpretation for an utterance a certain background premise would have to be used; and suppose that if this premise had been intended by the speaker, the fact that he had intended it to be used would be more relevant than the content of the utterance itself. This interpretation would automatically violate the principle of relevance because, by hypothesis, the proposition expressed by the utterance was not the most relevant proposition available to the speaker in the circumstances. Hence, this interpretation could not have been intended by a speaker attempting to observe the principle of relevance.

For example, imagine a student who, by breaking into the examiner's room at night, secretly knows that he has failed an exam. His professor, who has seen the results but is unaware of the break-in, says to him in casual conversation:

(32) Everyone who failed the exam will have his case considered at the next faculty meeting.

It is clear that by using the information in (33), the student can considerably increase the relevance of the utterance to him by deriving the contextual implication (34) rather than just (35):

(33) I have failed the exam.

(34) My case will be considered at the next faculty meeting.

(35) If I have failed the exam, my case will be considered at the next faculty meeting.

Indeed, he will certainly derive this implication. However, he will know that this could only be the *intended* interpretation if the professor had intended him to use (33) and derive the contextual implication (34). But if he *had* intended him to use it, he would have had to know that the student had secretly broken into the examiner's room, and the fact that he knew this would clearly be more relevant than (32) itself. The interpretation based on (33) could therefore not have been intended in these circumstances by a speaker observing the principle of relevance. Contrast this with the cases where speaker and hearer have already established (33) as part of their mutual knowledge and where, clearly, the speaker would intend the hearer to use (33) and derive (34).

Or take the case where an utterance has two interpretations, one with a normal degree of relevance and the other with considerably more than that. In the simplified model, the latter should be chosen. However, the full principle of relevance will (correctly) generally select the former. For instance, imagine two mothers chatting, and one saying to the other:

(36) My son has grown another foot.

This can mean either that her son has grown bigger or that he has grown an extra limb. The second interpretation is of course much more relevant. However, if this *were* the interpretation intended, then the principle of relevance would have been grossly violated. The speaker could indeed, if her son had become three-footed, produce a much more relevant utterance than a mere statement of the fact. She could elaborate in many ways while maintaining the level of relevance quite high, and not leave it to the hearer to work out all the crucial implications and raise all the questions that such an extraordinary fact would involve. Hence the full principle of relevance unquestionably selects the less dramatic and less relevant interpretations in such a case.

However, there are cases where the principle of relevance would not only fail to prevent mistakes, but would actually cause them to occur. If such mistakes do in fact occur, then the model of comprehension based on this principle is vindicated. We shall give two brief illustrations of these predicted mistakes.

We have argued that relevance is a function of the amount of processing and hence of the accessibility of the context required to derive contextual implications. Consider first the case of very touchy people who are easily hurt by the most innocent remark. Their behaviour and mistakes are easily predicted by the principle of relevance on the simple additional assumption that for them propositions about themselves and about what others may think of them constitute an inordinately rich and easily accessible potential context. It is an area which they are permanently attending to, and which therefore always provides a possible extension of the context. It will be appealed to and exploited whenever there is a chance. Such people will select the most relevant interpretation for them: that is, one whose main implications are about them. The checks that the principle of relevance provides in other cases are not likely to work here. Often, if the paranoid interpretation were correct, the speaker could, and hence should, have developed his remark and made it still more relevant. On the other hand, as the hearer will know,

there are social considerations and rules which may have prevented the speaker from being too explicit in his criticism.

Similarly, it is almost impossible to convince people who see a sexual intention in every other utterance that they are wrong: they select the most relevant interpretation for them, and easily account for its under-exploitation on the part of the speaker by considerations of propriety. Notice that if greater attention were paid to mutual knowledge, such mistakes would occur less.

A second example of a type of mistake predicted by the principle of relevance is the interpretation of the diviner's words by the consultant. Anthropologists have puzzled about why so many people keep trusting and consulting diviners. They have given a partial explanation by showing how the diviner's words were generally vague and open to a variety of interpretations, rather than downright wrong. It still remains to be explained why people should seek vague information from a diviner. The principle of relevance predicts that the consultant will automatically find the most relevant interpretation, and take it to be the one intended. Even though the consultant may have his doubts about the particular diviner he is consulting and want to test him (as is often reported to happen), it is enough that he should consider that divination is possible for the checks normally provided by the principle of relevance not to work. The diviner might know everything; hence there are no constraints at all on the premises he may have intended the consultant to use. He may have intended premises to be used that will only be available to the consultant at some time in the future. There is also a standard style of delivery and a standard degree of elaboration in the diviner's discourse which are taken to be part of the divinatory procedure. Hence the underexploitation on the part of the diviner of the interpretation arrived at by the consultant is no check on this interpretation.

IV. CONCLUSION

To return to the issue of mutual knowledge, we would argue that the best evidence for mutual knowledge (although it is by no means watertight) is not physical co-presence but rather comprehension: while we may construct very different descriptions of the same physical stimulus, adequate comprehension implies near identity of interpretation be-

tween interlocutors. Thus, if I understand what you say and give you no reason to doubt that I have understood it, then the interpretation of your utterance and all the contextual premises which were crucially used in arriving at it, can be assumed to be mutually known to us.

It takes very odd cases, such as those surveyed in Schiffer and Clark and Marshall (cf. *supra*) for a mistake in the assessment of mutual knowledge to cause a mistake in comprehension. This is not surprising if, as we have been arguing, mutual knowledge is inferred from comprehension rather than the other way round. On the contrary, mistakes in comprehension almost automatically cause mistakes in the assessment of mutual knowledge. Trivially, I will take my mistaken interpretation of your utterance to be part of common knowledge from the moment I have formed it.

Less trivially, the consultant assumes that the most relevant interpretation of the diviner's words, and all the premises used in arriving at it, were part of the diviner's knowledge from the start, and are mutually known from the moment he "understands" them. If consultants behaved as the mutual knowledge theorist would have them behave, they should assess mutual knowledge first. According to the mutual knowledge model, when they have doubts about the diviner they are consulting, and decide to test him, they should be easily able to see the vagueness of the predictions made, and should lose faith in each diviner they put to the test. Since diviners and divination have fared rather well, it is arguable that we should instead give up those models of comprehension based on an identification of context with mutual knowledge.

NOTES

1. We would like to thank Neil Smith for much valued encouragement and advice during the writing of this paper.
2. In an interesting paper, Hobbs (1979) has argued that a single general principle of coherence is involved in every aspect of interpretation. Although our proposals differ in detail, their aims are clearly similar.
3. This possibility is in fact the basis of Stoppard's play *Dogg's Hamlet Cahoot's Macbeth*, ed.