9

ON DEFINING RELEVANCE

DEIRDRE WILSON AND DAN SPERBER

1. *Introduction*

Under the category of RELATION, I place a single maxim, namely 'Be relevant'. Though the maxim itself is terse, its formulation conceals a number of problems that exercise me a good deal . . . I find the treatment of such questions exceedingly difficult, and I hope to revert to them in a later work. [Grice, 1975, p. 46.]

In these cautious terms Grice, in his William James Lectures, introduced a maxim of relevance which, together with maxims of truthfulness, informativeness, and perspicuity was to form the basis for his exciting and influential approach to pragmatics. In the years since then, in spite of the considerable amount of research done in and on the Gricean pragmatic framework, no explicit account of relevance suitable for use in pragmatic theory has emerged.[2] Indeed, there has been considerable scepticism over whether any such account is in principle possible.[3]

[1] We would like to thank the participants in the 1980 Pragmatics Workshop at Asilomar, California, and in particular Herb Clark, Paul Grice, Larry Horn, Jerry Morgan, Ellen Prince, Ivan Sag, and Bonnie Webber, for stimulating comments on an earlier version of this paper. We are also grateful to Diane Brockway, Robyn Carston, Jerry Katz, Ruth Kempson, and Yuji Nishiyama for their help.

[2] The following is a representative sample of comments:

'. . . no attempt to apply formal semantic theory to this notion [relevance] has been successful enough to provide a model that would be usable in pragmatics' [Thomason, 1973, p. 12].

'That relevance is relevant to linguistic description is painfully apparent . . . Equally apparent is the almost complete absence of any kind of formal linguistic treatment of the notion' [Gazdar, 1979, p. 45].

'. . . current accounts of conversational interaction depend crucially upon the undefined notion of "relevance"' [Werth, 1981, p. 130].

[3] See, for example, the following comments:

'Relevance is not a precise logical category . . . the word is used to convey an essentially vague idea' [Berlin, 1938–9, p. 21].

'Conventionally, considerations of relevance are apt to be relegated to the rhetorical rather than the logical dimension of assessment of arguments' [Haack, 1978, pp. 16–17].

Where definitions have been attempted, they have been approached in rhetorical, emotional, or aesthetic rather than logical terms.[4] Yet since utterance comprehension involves a substantial inferential element, it is not unreasonable to expect the drawing of inferences to contribute in some way to judgements of relevance. In this paper we want to argue that there is an intimate connection between inference and relevance, and that an approach to relevance in inferential terms can yield fruitful results for pragmatic theory.

2. *The nature of relevance*

It is usual to treat relevance as a property of utterances or a relation between an utterance and a text or discourse. However, relevant information may be derived not only from utterances and other communicative acts, but also from memory, observation, and inference. We shall treat relevance in the first instance as a property of propositions (information units, combinations of sense and reference); subsidiary definitions of relevance for utterances or discourses are relatively easy to construct. Similarly, we shall treat a proposition as relevant in the first instance, not to a text or discourse but to a context, where a context is a stock of information derived not only from preceding discourse, but also from memory, observation, and inference. From the formal point of view, a context is simply a set of propositions of arbitrary size and content. From the psychological point of view, both size and content are subject to considerable constraints; these will be ignored for the moment and discussed in a later section.

The interpretation of an utterance involves, on the one hand, identification of the proposition the speaker has expressed, and on the other hand, the processing of this proposition in a context provided by the hearer and consisting, as we have seen, of information derived from a variety of sources. To process a proposition is simply to extract

'. . . the notion of connection or dependence being appealed to here is too vague to be a formal concept of logic' [Suppes, 1957, p. 8].

'The difficulty of treating relevance with the same degree of mathematical sophistication and exactness characteristic of treatments of extensional logic led many influential philosopher-logicians to believe that it was *impossible* to find a satisfactory treatment of the topic' [Anderson and Belnap, 1975, p. xxi].

[4] For some recent attempts, see van Dijk, 1979, Dascal, 1977, and Werth, 1981. Van Dijk equates relevance with degree of importance in discourse, Dascal relates it to satisfaction of a conversational demand, and Werth equates it with appropriateness to the meaning of the previous utterance, together with the context.

information from it; to process it in a context is to supply additional background information which contributes in some way to its processing. The processing of a proposition may be treated in largely inferential terms. For example, the processing of (1) might involve the use of (1) as a premiss, together with the background information in (2), to deduce the conclusion in (3):

(1) Jackson has just bought a Rolls Royce, but his wife refuses to drive in expensive cars.
(2) The Rolls Royce is an expensive car.
(3) Jackson's wife refuses to drive in his Rolls Royce.

In this case, the processing of (1) involves going through the steps needed to draw a certain set of inferences, and the context will be, or at least include, the set of background assumptions, such as (2), used as supplementary premisses in the inferential process.

When a proposition P is added to a context $C_1 \ldots C_n$, two distinct types of inference process involving P as a premiss may take place. On the one hand, P may be taken as the sole premiss, and its logical implications obtained. These are, of course, invariant from context to context, but we can say that P is *informative* in the context $C_1 \ldots C_n$ iff P has at least one logical implication not implied by $C_1 \ldots C_n$.

The processing of P in a context $C_1 \ldots C_n$ may also yield a further set of implications, this time context-dependent: the set of propositions which are logically implied not by P alone, nor by $C_1 \ldots C_n$ alone, but by the union of P and $C_1 \ldots C_n$. Call these the *contextual implications* of P in the context $C_1 \ldots C_n$. For example, consider a context consisting of the two propositions (4a) and (4b):

(4) a. If the Chairman resigns, Jackson will take over his duties.
b. If Jackson takes over the Chairman's duties, the company will go bankrupt.

If proposition (5) is added to this context, conclusions (6) and (7) may be obtained as contextual implications of (5) in the context (4a)-(4b):

(5) The Chairman has resigned.
(6) Jackson will take over the Chairman's duties.
(7) The company will go bankrupt.

Propositions (6) and (7) are contextual implications of (5) in this context: they are logically implied by the union of (4) and (5), but by neither (4) alone nor (5) alone.

As informativeness in a context is definable in terms of logical implication, relevance in a context might be approached in terms of contextual implication. We shall say that a proposition is relevant in a context $C_1 \ldots C_n$ iff it has at least one contextual implication in $C_1 \ldots C_n$. Intuitively, being relevant in a context is a matter of connecting up with the context in some way. According to this proposal, relevance in a context is a matter of connecting up with the context in a highly specific way: so as to have contextual implications in that context. By this definition, because (5) has contextual implications in the context (4), it would be relevant in that context as well.[5]

This inferential approach to relevance would fit naturally into an inferential theory of comprehension. In processing a proposition in a context, the hearer would automatically derive its contextual implications, on the basis of which the relevance of the proposition would be established. In fact, one might go further and claim that the purpose of processing a proposition is precisely to establish its relevance, making relevance the foundation stone of pragmatic theory. Whatever its initial plausibility, however, this proposal can not be seriously entertained without answers to a number of questions being provided. Here we shall briefly consider the three that seem to us most urgent. It is perhaps worth mentioning that none is exclusive to our particular approach. The first, which has to do with the nature of the inference rules used, arises in any inferential pragmatic theory. The second, which has to do with the fact that relevance is a matter of degree, arises in any theory in which relevance plays a role. The third, which has to do with the nature and identification of contexts, arises in any pragmatic theory in which context plays a role. The answers we shall give, unlike the questions themselves, are often possible only within the sort of framework that we propose.

3. *The deductive system*

Suppose the logic used in deriving contextual implications is a standard one, say a standard natural deduction system. Then for any context Q and any proposition P, the conjunction of P and Q will be a contextual implication of P in the context Q: *P and Q* is logically implied by the union of P and Q, but by neither P alone nor Q alone. Every proposition will have at least one contextual implication in every context and hence, according to our definition of relevance, will be relevant in

[5] Goodman 1961 defines a notion of 'relative aboutness' along somewhat similar lines, although with rather different purposes.

every context; which is absurd. Therefore, either the logic used in deriving contextual implications is not a standard one, or our approach to relevance is incorrect.

In fact, quite independently of this approach to relevance, there is good reason to think that the logic used in utterance comprehension is not a standard one. On the one hand, it must be much more extensive, providing rules for every concept that can play a role in the inferential processing of propositions, including many that are of no particular interest to logicians. On the other hand, it must be more restrictive in certain ways. For example, the standard logical implications of a single proposition P include many propositions which would never in fact be derived during comprehension of an utterance expressing P:

(8) a. $P \& P$
b. $P \vee P$
c. $--P$
d. $P \vee Q$
e. $-P \rightarrow Q$
f. $Q \rightarrow P$.

It is in general simply false that a hearer, given an utterance expressing P, might think that any of the implications in (8) was part of the speaker's intended message, and any inferential account of comprehension must provide some method of excluding them.

The need for some restriction is compounded by the fact that the rules which give rise to (8) may reapply to their own output, yielding infinite sets of logical implications along the lines in (9):

(9) a. $(P \& P) \& P \ldots$
b. $(P \vee P) \vee P \ldots$
c. $\ldots ----P$
d. $(P \vee Q) \vee R \ldots$
e. $\ldots -(-P \rightarrow Q) \rightarrow R \ldots$
f. $\ldots R \rightarrow (Q \rightarrow P)$.

The full set of logical implications of a proposition could thus not in principle be drawn in any step-by-step way during the finite amount of time it takes to process a proposition.

The inferences in (8) and (9) are in some sense trivial. The rules which give rise to them are fairly easy to characterize: they are rules which may apply to any proposition at all, regardless of its form or content, and which may thus reapply indefinitely to their own output.

There are two possible reactions to the problem of trivial inferences and the rules which give rise to them. Most theories of comprehension which consider the problem at all retain something equivalent to a standard logic, but attempt to restrict its functioning in some way. It is felt that in certain circumstances a trivial rule might be needed, and it would be too drastic a move to eliminate all trivial rules entirely.[6] We feel, on the contrary, that the more drastic approach is correct.

Our hypothesis is that the deductive system used in spontaneous information processing is a purely interpretive one, in the sense that each of its rules requires, for its application, the presence of a particular concept in the proposition or propositions being processed. In other words, each rule is essentially an elimination rule. Such a system could contain something like the standard rule of *and*-elimination, which applies only to propositions in which *and* is present, but could have no equivalent of the standard rule of *and*-introduction, which imposes no conditions on the form or content of the propositions to which it applies. Such a system, though it could contain elimination rules for a wide range of concepts not treated at all in standard logics, could contain no rule which permits the derivation of any of the propositions in (8) or (9) from the single premiss P.

In a system of this type, the problem of trivial inferences would not arise. In particular, since there would be no rule of *and*-introduction, it would not be possible, given a proposition P and an arbitrary context Q, to derive the trivial implication $P \,\&\, Q$ as a contextual implication of P in the context Q. The elimination of the problem of trivial inference thus automatically provides a solution to the first of the problems raised by the proposed account of relevance.[7]

The inferential processing of a proposition could now be conceived of as carried out by an automaton which, given a proposition, a finite

[6] See, for example, Johnson-Laird, 1975.

[7] This is the only one of our three problems to be treated by relevance logicians. The system we envisage has some similarity to Parry's system of 'analytical implications' as outlined in Anderson and Belnap, 1975. It is sometimes suggested that a rule of *and*-introduction would be needed to derive R in the following circumstances: the context contains P and $(P \,\&\, Q) \rightarrow R$, and the new information to be processed is Q. However, we are not claiming that the most empirically adequate system of elimination rules contains only those envisaged in standard logics. There is good reason to think that in the above circumstances, the hearer would automatically convert $(P \,\&\, Q) \rightarrow R$, in a single step, to $P \rightarrow (Q \rightarrow R)$, which would combine with P to yield $Q \rightarrow R$ before Q was ever presented. On presentation of Q, the hearer could thus proceed directly to the conclusion R. For further discussion, see 'Reply to Gazdar and Good', in Sperber and Wilson, 1982.

context, and a set of non-trivial inference rules as input, would derive the full, finite set of non-trivial logical implications of the union of the proposition with the context as output. These would include the contextual implications of the proposition in the context. One could then compare the contextual implications that a given proposition would have in different contexts, or that different propositions would have in the same context. One could also compare the *amount* of processing that different propositions would require in a given context, or that a given proposition would require in different contexts, where the amount of processing is the number of steps that some automaton would have to go through in order to be sure of deriving all the contextual implications of a proposition in a context. Leaving special cases and technicalities aside, amount of processing is determined by, on the one hand, the number of non-trivial logical implications of the proposition being processed, and on the other hand, the number of non-trivial logical implications of the context. The simplest effective automaton would merely examine each of the members of the cartesian product of these two sets to check whether any inference rule applied. Roughly speaking, the greater the semantic complexity of the proposition being processed, and the larger the context, the greater the amount of processing that will be required.[8] We shall argue that degrees or comparisons of relevance are based on assessments of numbers of contextual implications on the one hand and amount of processing on the other.

4. *Degrees of relevance*

Grice's maxim simply says 'Be relevant'. If, as we have suggested, being relevant is merely a matter of expressing a proposition which has at least one contextual implication in a context accessible to the hearer, such a maxim would constrain the speaker, and hence guide the hearer, hardly at all. However, if degrees of relevance could be defined, or some basis for comparisons of relevance could be given, then some much more constraining maxim such as 'Try to be as relevant as is possible in the circumstances' might be proposed, and interpretation would take place under correspondingly severer constraints. In this section, we shall provide a basis for comparing the relevance of different propositions in

[8] There is, of course, a third factor affecting the amount of processing that a certain proposition requires in a given context, namely the number of contextual implications itself: each contextual implication adds one derivational step. However, we can ignore this factor, since it does not contribute to an assessment of amount of processing independent of the assessment of number of contextual implications, which is what we are after.

a fixed context, going on in the next section to deal with context selection and assessment in variable contexts. The examples considered will be highly unrealistic, partly because utterance comprehension does not take place in fixed contexts, and partly because actual contexts are unlikely to be as small as those used here. We shall also make the simplifying assumption that speakers not only aim at maximal relevance, but succeed in their aim, so that the hearer's task is merely to choose the interpretation that is maximally relevant. In a fuller account this simplifying assumption would be dropped. Examples and discussion should thus be taken only as an illustration of the criteria around which a fuller account could be constructed.

If relevance is linked to contextual implications, it seems reasonable that the more contextual implications a proposition has in a given context, the more relevant it is. We shall say that, other things being equal, the more relevant of two propositions in a given context will be the one with more contextual implications. However, a further factor must be taken into account. Two propositions may have the same number of contextual implications in a certain context, but one may be more semantically complex than the other, and contain whole stretches of information that do not connect up with the context and make no contribution to relevance. Intuitively, the existence of this extraneous information detracts from the relevance of the proposition. Because it also increases the amount of processing required, we shall say, for reasons that will become clearer in the next section, that, other things being equal, the more relevant of two propositions in a given context is the one which requires least processing. In assessments of relevance, there are thus two factors to take into account: on the one hand, numbers of contextual implications, and on the other hand, the amount of processing needed to obtain them.

To illustrate this, take the context in (10):

(10) a. The tickets cost £1 each.
b. People may buy more than one ticket.
c. The person who bought the green ticket wins the prize.
d. The prize is £100,000.
e. Anyone who wins £100,000 can fulfil the dreams of a lifetime.

If a hearer had only the information in (10), and a speaker knew that all of (11)-(13) were true, which of these propositions should the speaker, aiming at maximal relevance, choose to express?

(11) James bought the green ticket.
(12) Charles bought the blue ticket.
(13) James bought the green ticket, and today is Tuesday.

(11) and (12) require roughly the same amount of processing,[9] since they have parallel logical structures and are being processed in the same context. Of these two propositions, it is therefore the one with the most contextual implications that will be the more relevant. While (12) has only the single contextual implication (14), (11) has the set of contextual implications (15)-(18):

(14) The blue ticket cost Charles £1.
(15) The green ticket cost James £1.
(16) James wins the prize.
(17) James wins £100,000.
(18) James can fulfil the dreams of a lifetime.

Hence (11) is more relevant than (12) in this context.

In comparing (11) and (13), the same two factors—number of contextual implications and amount of processing—must be taken into account. In this case, (11) and (13) have exactly the same contextual implications, but require different amounts of processing. Because (13) entails (11), (13) requires every step of processing that (11) requires, and more besides. By the proposed criteria, it is the proposition which requires the smallest amount of processing—that is, (11)—that is the more relevant in the context. Of the three propositions (11)-(13), it is thus (11) that is selected as most relevant by the proposed criteria, and this accords with intuitive judgements of relevance.

These criteria do not, and need not, always yield such clear-cut results. For example, they do not say what the result would be of comparing (12) and (13), which differ both in the amount of processing required and in their numbers of contextual implications. They are designed to yield clear-cut results in just those cases where human beings can make clear-cut judgements of relative relevance. The sort of comparisons among different propositions that a hearer is called on to make during utterance comprehension are generally more like those in (11) and (12), and (11) and (13), than (12) and (13). For example, a hearer may have to choose among various candidate disambiguations and reference assignments for a given utterance. In most cases of

[9] With the unimportant qualification suggested in footnote 8.

reference assignment, and many of disambiguation, the candidate propositions to be compared will share much of their logical structure, as do (11) and (12), and may enter into entailment relations, as do (11) and (13). Moreover, in the normal case the alternative candidates differ so grossly in numbers of contextual implications that comparisons are quite straightforward, and only a single proposition will be worth considering. When no single candidate clearly emerges, we predict that the hearer will be unable to decide which interpretation is intended.[10]

So far we have assumed that contexts are arbitrary sets of propositions, fixed in advance. In looking at examples (11)-(13) above, the reader may have been tempted to expand context (10) in one way or another to increase the relevance of the proposition being processed: introducing, for instance, assumptions which would combine with the extraneous information in (13) to permit further contextual implications to be derived. In real life, contexts are not fixed in advance, but are chosen at least partly in function of the proposition being processed. It is to the question of context selection, or context construction, that we now turn.

5. *Context selection*

The problem of context selection is not unique to our framework. Any adequate theory of comprehension must describe the role of background assumptions in utterance interpretation, and the principles by which they are selected to play this role. In our framework, the role of background assumptions is to provide premisses which will combine with the proposition being processed to yield contextual implications. The goal of processing is in general to maximize, as far as possible, the relevance of the proposition being processed: that is, to obtain the maximum of contextual implications in return for any given amount of processing. We want to show that because variations in context may increase or decrease the relevance of the proposition being processed, the goal of maximizing relevance may simultaneously guide the choice of context.

The hearer of an utterance has available a set of potential contexts from which an actual context must be chosen. We assume that there is

[10] In a fuller account, it could be shown that a speaker who did not believe that one clear candidate would emerge, or who caused the hearer to hesitate between alternative interpretations, would not have been observing the principle of maximal relevance. See Sperber and Wilson, 1986, for further discussion.

a small, immediately accessible context, fixed in advance, and consisting of the proposition which has most recently been processed, together with its contextual implications. When new information is received, say from an utterance, it will be processed in this most immediate context. If the initial context is (19) and the proposition expressed by the utterance is (20), some degree of relevance is immediately achieved:

(19) If the interest rate has risen, the company will go bankrupt.
(20) The interest rate has risen.

It may happen, however, that unless the initial context is extended in some way, no degree of relevance can be achieved. This would be so if, with the same initial context, the proposition expressed were not (20) but one of (21)-(23):

(21) The interest rate has done what you said it would do.
(22) The interest rate has done what the sun does every morning.
(23) The interest rate has done this. [speaker demonstrates]

If the goal of processing is to maximize the relevance of the information being processed, the hearer will be forced to add to the initial context (19) further information, which may be remembered from earlier discourse (in the case of (21)), recovered from encyclopaedic memory (in the case of (22)), or derived from sense perception (in the case of (23)). The goal will be to find, from the most immediately accessible source, premisses which will combine with the proposition beng processed, to yield the maximum of contextual implications in return for the available amount of processing.

As these examples show, the accessibility of potential contexts may be altered by the content of the proposition being processed, which may direct the hearer's attention in one case to the physical environment, in another to the preceding discourse, in another to encyclopaedic memory. More complex examples require simultaneous or sequential extensions of the context in a variety of different directions, and there is no principled limit on the number of extensions that may be needed to establish the relevance of a given proposition. Each extension provides a new potential context in which the proposition could be processed, up to a maximal context consisting of the entire contents of the hearer's accessible memory.

If there were only one direction of extension, the set of potential contexts for processing a given proposition would be strictly ordered, ranging from a smallest, most easily accessible initial context, through

ever larger, less accessible extensions, up to a most inclusive, least accessible context containing everything the hearer knows. Because of the variety of possible directions of extension, the actual situation is rather more complex; however, there is a partial ordering of potential contexts, with a series of ever more inclusive contexts extending out from the initial context in various directions.

If the set of potential contexts is structured in this way, every extension of the initial context will incur a double cost in processing terms. First, each extension increases the size of the context, and the larger the context, the greater the amount of processing. Second, later extensions are harder to access, and accessing itself is presumably a step-by-step procedure with associated costs in processing. Hence, every extension of the context increases the cost of processing, and must be expected, other things being equal, to decrease the relevance of the proposition being processed.

We shall say that a proposition is relevant in a set of contexts $K_1 \ldots K_n$ iff it is relevant in at least one context which is a member of that set, and that it is more relevant in K_1 than K_2 if, other things being equal, it has more contextual implications in K_1 than in K_2, or, other things being equal, the amount of processing required to obtain the implications it has in K_1 is smaller than the amount required to obtain the implications it has in K_2. Maximizing the relevance of a given proposition is therefore a matter of choosing a context which maximizes its contextual implications and minimizes the amount of processing; in other words, it is a matter of maximally efficient processing.

This approach to the assessment of relevance in variable context can be illustrated with the following highly simplified example. Assume that there is an initial context (24C1), to which (24C2), (24C3) and (24C4) may be added, in that order:

(24) C1: Jackson has chosen the date of the meeting.

C2: If the date of the meeting is February 1st, the Chairman will be unable to attend.

C3: If the Chairman is unable to attend, Jackson's proposals will be accepted.

C4: If Jackson's proposals are accepted, the company will go bankrupt.

A hearer aiming to maximize the relevance of (25) or (26) in this set of contexts should only be willing to incur the extra costs of extending

the context in return for a compensating increase in the number of contextual implications:

(25) The date of the meeting is February 1st.
(26) The date of the meeting is February 5th.

In the initial context C1, (25) has the single contextual implication (27), and (26) the single contextual implication (28):

(27) Jackson has chosen February 1st as the date of the meeting.
(28) Jackson has chosen February 5th as the date of the meeting.

In this context, both propositions are relevant, and indeed, because of their similarities in logical structure, context, and contextual implications, both will be equally relevant.

If the initial context is extended to include C2, (25) gains the further contextual implication (29):

(29) The Chairman will be unable to attend.

The extension would be worthwhile if the cost of accessing C2 was not so great as to outweigh the gain in contextual implications. If the context is further extended to include C3, and then C4, the additional contextual implications (30) and (31) are in turn obtained:

(30) Jackson's proposals will be accepted.
(31) The company will go bankrupt.

Again, if the accessing costs are not too high, these extensions would both be worthwhile. Thus, someone processing (25) in the set of contexts C1-C4 might have good reason to continue expanding the context up to its maximum size, because every expansion would increase the relevance of (25).

In the case of (26), the results of any extension beyond the initial context C1 would be quite different. Each extension incurs an extra cost in processing, but without any compensating increase in contextual implications. Whichever context is chosen, only the single contextual implication (28) is obtained. Thus, someone processing (26) in the set of contexts C1-C4 would have no reason to extend the context beyond the initial stage C1, because every expansion would decrease the relevance of (26).

Within this inferential framework, it is thus possible to see in principle how context selection might take place, and more specifically, how a proposition might help to determine its own context, subject to

constraints of accessibility and relevance. In a more realistic account, more serious attention would have to be given to the choice of particular directions for extension of the context. For example, the hearer of (26) would probably not give up at context C1, but would look for relevance in another direction. Similarly, more attention would have to be paid to the question of how processing costs are balanced against numbers of contextual implications in real-life situations. For example, at what point would the hearer of (25) decide that the cost of new extensions to the context had become too high? These are empirical questions, which we have not tried to solve directly, and to which, moreover, we believe that there will be different answers for different people in different circumstances.

The relation of our work to these and other such empirical questions in the theory of comprehension is indirect. For example, we do not believe that a hearer, in disambiguating an utterance, will actually compare all its possible interpretations and rank them for relevance before deciding what has been said. What we have tried to describe is not the procedures used in disambiguation but the goal these procedures are designed to achieve, the property they are designed to diagnose. Our claim is that the interpretation the speaker should have intended, and the one the hearer should choose, is the one that satisfies a principle of maximal relevance. Knowing the goal, it should become easier to describe the procedures, and it is in this way that our work may contribute to descriptive work on comprehension.

The assumptions crucial to our framework are, first, that comprehension is a largely inferential process, the role of context being essentially to provide premisses for the calculation of contextual implications, and second, that the contextual implications of a proposition in a finite context are themselves finite. With these assumptions it is possible to see how context selection might in principle take place; without them the framework would collapse. What is in some ways less important is our choice of terminology. For us, 'relevance' is a technical term, designed to play a role in an overall pragmatic theory rather than to approximate everyday usage. We do believe that a quite substantial range of everyday intuitions about relevance are captured within this theory, and hence that this choice of terminology is quite appropriate. However, the terminology could be questioned without the overall framework being affected at all.

6. *Conclusion*

Grice regards the practice of observing the maxims of conversation as not merely 'something that all or most do IN FACT follow but as something that it is REASONABLE for us to follow, that we SHOULD NOT abandon' [Grice 1975, p. 48]. He is uncertain why this should be so, and suggests that an answer must await a clarification of 'the nature of relevance and the circumstances in which it is required' [p. 49]. The answer suggested by our attempted clarification is as follows.

In cognitive terms, what human beings are looking for in the information they process is relevance. Processing involves effort, and will only be undertaken in expectation of some reward in terms of contextual implications; the greater the expected reward, the more effort one will be prepared to undertake. Expectations of relevance are not constant across all individuals and circumstances, and a speaker or writer who consistently disappoints the expectations of an audience will cause a downward readjustment in their subsequent expectations, and may lose their attention entirely. If we want you to read our work, it is thus in our interest to aim at as high a level of relevance as we possibly can, creating expectations which will induce you to continue reading our work with the attention we think it deserves. Given the brute facts of cognitive psychology, this is rational, or at least reasonable, behaviour, although not particularly altruistic or virtuous.

It follows that information that has been deliberately communicated, unlike information from other sources, comes with a guarantee that a certain level of relevance has been attempted, if not achieved. It is this guarantee that makes the difference between merely uttering sentences in front of someone and saying something to someone, that justifies the disambiguation of meaning and the calculation of implicatures, that defines a separate field of pragmatics. Pragmatics so conceived is about the way in which the universal search for relevance is both served and exploited in verbal communication, and the study of relevance is its foundation stone.[11]

REFERENCES

Anderson, A. R. and N. D. Belnap, *Entailment*, Princeton University Press, 1975.

[11] For discussion of Grice's contribution to pragmatics, see Wilson and Sperber, 1981; for further discussion of the framework described in this paper, see Sperber and Wilson, 1982; for an attempted outline of a pragmatic theory, see Sperber and Wilson, 1986.

Berlin, I., 'Verification', in *Proceedings of the Aristotelian Society*, 1938–9.

Dascal, M., 'Conversational Relevance', in A. Margalit (ed.), *Meaning and Use*, D. Reidel, 1979.

van Dijk, T., 'Relevance Assignment in Discourse Comprehension', in *Discourse Processes*, 2 (1979), pp. 113–126.

Gazdar, G., *Pragmatics: Implicature, Presupposition and Logical Form*, Academic Press, 1979.

Goodman, N., 'About', in *Mind*, LXX, 277 (1961), pp. 1–24.

Grice, H. P., 'Logic and Conversation', in P. Cole and J. Morgan (eds), *Syntax and Semantics 3: Speech Acts*, Academic Press, 1975, pp. 41–58.

Haack, S., *Philosophy of Logics*, Cambridge University Press, 1979.

Johnson-Laird, P. N., 'Models of Deduction', in R. J. Falmagne (ed.), *Reasoning: Representation and Process*, Lawrence Erlbaum, Hillsdale, NJ, 1975.

Sperber, D. and D. Wilson, 'Mutual Knowledge and Relevance in Theories of Comprehension', in N. V. Smith (ed.), *Mutual Knowledge*, Academic Press, 1982, pp. 61–131.

—— and ——, *Relevance: Communication and Cognition*, Blackwell, 1986.

Suppes, P., *Introduction to Logic*, D. van Nostrand, 1957.

Thomason, R., 'Semantics, Pragmatics, Conversation and Presupposition', unpublished mimeo, Philosophy Dept. University of Pittsburgh, 1973.

Werth, P. N., 'The Concept of Relevance in Conversational Analysis', in P. N. Werth (ed.), *Conversation and Discourse*, Croom Helm, 1981, pp. 129–54.

Wilson, D. and D. Sperber, 'On Grice's Theory of Conversation', in P. N. Werth (ed.), *Conversation and Discourse*, Croom Helm, 1981, pp. 155–78.