

Meaning  
and  
Interpretation

EDITED BY

*Charles Travis*

Basil Blackwell

B040  
.M45  
1986

# Contents

© Basil Blackwell 1986

First published 1986

Basil Blackwell Ltd  
108 Cowley Road, Oxford OX4 1JF, UK

Basil Blackwell Inc.  
432 Park Avenue South, Suite 1505,  
New York, NY 10016, USA

All rights reserved. Except for the quotation of short passages for the purposes of criticism and review, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

## *British Library Cataloguing in Publication Data*

Meaning and interpretation.  
1. Semantics (Philosophy)  
I. Title  
149'.946 B820

ISBN 0-631-14644-X

## *Library of Congress Cataloging in Publication Data*

Meaning and interpretation.  
Bibliography: p.  
Includes index.  
1. Semantics (Philosophy)—Addresses, essays, lectures. 2. Meaning (Philosophy)—Addresses, essays, lectures. I. Travis, Charles, 1943—  
B840.M45 1986 121'.68 85-26708  
ISBN 0-631-14644-X

Typeset by Photographics, Honiton, Devon  
Printed in Great Britain by Page Bros, Norwich

Preface vii  
Introduction *Charles Travis* 1

## Part I Words in Language and Use

- 1 Roles and connecting paths  
*Gilles Fauconnier* 19
- 2 Inference and implicature  
*Deirdre Wilson and Dan Sperber* 45
- 3 Ambiguity and the semantics-pragmatics distinction  
*Ruth Kempson* 77
- 4 Domains of discourse and common-sense metaphysics  
*Adam Morton* 105

## Part II Structure in Language and Thought

- 5 Tacit knowledge, and the structure of thought and language  
*Martin Davies* 127
- 6 Conceptual structure  
*John Campbell* 159

## Part III Content and Truth

- 7 Meaning and metaphysics  
*S. D. Guttenplan* 177
- 8 How can we tell whether a commitment has a truth condition?  
*Simon Blackburn* 201



path) which endows the expression with a certain number of potential interpretations.

The overall view is that an expression specifies some construction to be performed, but that this construction may be underdetermined with respect to the space configuration already set up when the expression occurs: there may be several ways to perform the appropriate construction. For example, the definite description 'head of state' in (44) instructs us to find an element with the property 'head of state', and to identify another element linked to the first by an appropriate connecting path. The number of possibilities will depend on how many spaces we have containing an element 'head of state', and on the length of the connecting paths originating from these elements. There is no theoretical limit on this number but, of course, in each specific configuration it will be finite and even usually very small. The number of readings cannot be a property of the expression *per se*: it results from the possible combinations of the instructions with space configurations.

The standard attribution of ambiguity to expressions is done in isolation, by constructing a minimal context and therefore a minimum space configuration. Since this procedure is standard and reproducible, it creates the illusion that the number of readings is a property of the expression itself. But, in fact, it is a property of the expression in minimal contexts and does not project to the general case.

## 2

### Inference and implicature

*Deirdre Wilson and Dan Sperber*

Recent accounts of utterance interpretation have tended to downplay the role of deductive reasoning in comprehension. Here are two illustrations from current textbooks on discourse analysis. Brown and Yule say: 'It may be the case that we are capable of drawing a specific conclusion . . . from specific premises . . . via deductive inference, but we are rarely asked to do so in the everyday discourse we encounter. . . . We are more likely to operate with a rather loose form of inferencing.'<sup>1</sup> Similar views are expressed by de Beaugrande and Dressler: 'Humans are evidently capable of intricate reasoning processes that traditional logics cannot explain: jumping to conclusions, pursuing subjective analogies, and even reasoning in the absence of knowledge. . . . The important standard here is not that such a procedure is logically unsound, but rather that the procedures work well enough in everyday affairs.'<sup>2</sup>

The implication of these remarks is that any element of deductive reasoning in the 'loose' or 'unsound' inference processes involved in utterance interpretation can safely be ignored. Our own researches have led us to a quite different conclusion. We do believe that non-demonstrative inference plays a crucial role in utterance interpretation. We also believe, however, that deductive

We are grateful to Rose Maclaran for a number of helpful comments on an earlier version.

1 G. Brown and G. Yule, *Discourse Analysis* (CUP, Cambridge, 1983), p. 33.

2 R. de Beaugrande and W. Dressler, *Introduction to Text Linguistics* (Longman, London, 1981), pp. 93-4.

reasoning makes a crucial contribution to the non-demonstrative inference processes spontaneously used in utterance interpretation. In this chapter we will offer a partial justification of these claims by illustrating the contribution of deductive reasoning to just one aspect of utterance interpretation in which non-demonstrative inference is generally agreed to play a major role: the recovery of implicatures. A fuller justification would deal with other aspects of interpretation – disambiguation, reference assignment, the recovery of illocutionary force and the interpretation of figurative language, for example – and would not merely illustrate the contribution of deductive reasoning to spontaneous non-demonstrative inference processes but would provide a psychologically adequate account of these non-demonstrative inference processes themselves.<sup>3</sup>

#### THE CALCULABILITY REQUIREMENT ON IMPLICATURES

Grice distinguishes two main types of implicit content or implicature that an utterance can convey.<sup>4</sup> *Conventional implicatures* are determined by particular lexical items or linguistic constructions occurring in the utterance. *Conversational implicatures* follow from general maxims of truthfulness, informativeness, relevance and clarity that speakers are assumed to observe. For Grice, the important difference between conventional and conversational implicature is that the conventional implicatures of an utterance are arbitrarily stipulated, whereas its conversational implicatures should be recoverable by a reasoning process.<sup>5</sup> The presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a *conversational implicature*: it will be a *conventional implicature*.<sup>5</sup> Grice regards this calculability requirement as fun-

<sup>3</sup> This fuller justification is attempted in D. Sperber and D. Wilson, *Relevance: Communication and Cognition* (Basil Blackwell, Oxford, 1986).

<sup>4</sup> H. P. Grice, 'Logic and conversation', in *Syntax and Semantics 3: Speech Acts*, eds P. Cole and J. Morgan (Academic Press, New York, 1975), pp. 41–58.

<sup>5</sup> *Ibid.*, p. 50.

damental: 'The final test for the presence of a conversational implicature had to be, as far as I could see, a derivation of it. One has to produce an account of how it could have arisen and why it is there. And I am very much opposed to any kind of sloppy use of this philosophical tool in which this condition is not fulfilled.'<sup>6</sup> However, his own account of the derivation process is rather sketchy, and although the idea of conversational implicature has had enormous appeal and been used in an informal way to account for a wide range of pragmatic phenomena, little progress has been made in specifying the exact nature of the inference process by which conversational implicatures are 'worked out'.

Grice suggests that the conversational implicatures of an utterance might be derived by arguments of the following form:<sup>7</sup>

- (1) (a) He has said that *p*.
- (b) There is no reason to suppose that he is not observing the maxims.
- (c) He could not be doing this unless he thought that *q*.
- (d) He knows (and knows that I know that he knows) that I can see that the supposition that he thinks that *q* is required.
- (e) He had done nothing to stop me thinking that *q*.
- (f) He intends me to think, or is at least willing to allow me to think, that *q*.
- (g) And so, he has implicated that *q*.

It is unclear what sort of argument this is meant to be; it is not even clear which of (1a–g) are meant to be premises and which conclusions. What does seem clear is that (1c), in which the content of the implicature is introduced for the first time, is not directly deducible from (1a–b). Either (1c) is simply an independent premise, or it is meant to be derivable from (1a–b) with the aid of some supplementary premises whose nature has been left unspecified. What (1) really offers is not a method for working out

<sup>6</sup> H. P. Grice, 'Presupposition and conversational implicature', in *Radical Pragmatics*, ed. P. Cole (Academic Press, New York, 1981), pp. 183–98, at p. 187.

<sup>7</sup> Grice, 'Logic and conversation', p. 50.

the content of the propositions that the speaker, in producing an utterance, implicitly commits himself to, but rather a method for working out which of these commitments the speaker *meant*, in Grice's special technical sense.<sup>8</sup> An adequate pragmatic theory should also provide some method of recovering the content of the implicatures themselves: that is, some method of deriving not (1g) but (1c).

It is becoming a commonplace of the pragmatic literature that deductive inference plays little if any role in the recovery of implicatures. Leech says that implicatures are 'probabilistic', and that the process by which they are recovered 'is not a formalized deductive logic, but an informal rational problem-solving strategy'.<sup>9</sup> Levinson says that implicatures 'appear to be quite unlike logical inferences, and cannot be directly modelled in terms of some semantic relation like entailment'.<sup>10</sup> Bach and Harnish say that the form of inference by which implicatures are recovered 'is not deductive but what might be called an inference to a plausible explanation'.<sup>11</sup> Brown and Yule, who as we saw above claim that utterance comprehension rarely involves deductive processes, add that in the recovery of implicatures 'we are more likely to operate with a rather loose form of inferencing'.<sup>12</sup> These remarks can be taken in a number of ways.

In a sense, they follow directly from Grice's characterization of implicature, and some of these authors may be making a purely definitional point. Consider (2b), for example:

- (2) (a) *He*: Will you have some coffee?  
(b) *She*: Coffee would keep me awake.

In normal circumstances, the speaker of (2b) would implicate (3):

- (3) She won't have any coffee.

8 For discussion, see H. P. Grice, 'Utterer's meaning, sentence meaning and word meaning', *Foundations of Language*, 4 (1968), pp. 1-18.

9 G. N. Leech, *Principles of Pragmatics* (Longman, London, 1983), pp. 30-1.

10 S. C. Levinson, *Pragmatics* (CUP, Cambridge, 1983), pp. 115-16.

11 K. Bach and R. M. Harnish, *Linguistic Communication and Speech Acts* (MIT Press, Cambridge, Mass., 1979), pp. 92-3.

12 Brown and Yule, *Discourse Analysis*, p. 33.

Now (3) is not deducible from the content of (2b) alone: (4) is not a contradiction:

- (4) Coffee would keep her awake, and she will have some coffee.

Indeed, if (3) *were* deducible from (2b) it would not be an implicature in Grice's sense, since according to him 'the truth of a conversational implicature is not required by the truth of what is said (what is said may be true - what is implicated may be false)'.<sup>13</sup> If this is taken as a defining feature of implicature, no implicature will be deducible from the explicit content of an utterance alone.

However, to show that (3) is not directly deducible from (2b) is not to show that deduction plays no significant role in its derivation. Grice himself claims that background knowledge must play a role in the process by which conversational implicatures are 'worked out'. Why should the hearer not simply supply the background assumptions in (5) and use them, together with the content of (2b), to deduce the conclusion in (3)?

- (5) (a) She doesn't want to be kept awake.  
(b) She won't have anything that would keep her awake.

We will argue that deduction processes of this type play a central role in the recovery of implicatures. This is not, of course, to tell the whole story about how implicatures are derived. It is also necessary to show how appropriate premises for the deduction process are selected and potential conclusions evaluated: why does the hearer of (2b) supply the background assumptions (5) and accept the conclusion (3) rather than, say, supplying the assumptions (6) and deriving the conclusion (7)?

- (6) (a) She wants to be kept awake.  
(b) She will have anything on offer that would keep her awake.  
(7) She will have some coffee.

However, there is no reason in principle why an account of implicature in which deduction is a central element should not be

13 Grice, 'Logic and conversation', p. 58.

adequate to deal with the interpretation of (2b) and other examples of implicature-carrying utterances.

Some of the authors cited above seem to be making not just the definitional point that implicatures cannot be deduced from the content of the utterance alone, but the stronger claim that deduction plays no significant role in their derivation, and in particular that an account like the one just suggested cannot be correct.<sup>14</sup> To be able to evaluate this claim we would have to have not only a clearer idea of the deduction-based account, but also some idea of possible alternatives to it. By what non-deductive inference processes might the hearer of (2b) arrive at the conclusion in (3), and how would the alternative conclusion in (7) be ruled out?

It is now fairly widely recognized that there can be no non-demonstrative inference *rules* (in the sense that there are deductive inference rules) which, given a set of premises, simply enumerate a set of valid conclusions. Instead, the process of reaching valid non-demonstrative conclusions is standardly broken down into two distinct stages: hypothesis formation and hypothesis confirmation. For example, the hearer of (2b) would have first to form the hypothesis in (3), or the hypothesis that the speaker was trying to communicate (3), and second to confirm or disconfirm this hypothesis.

As Fodor points out,<sup>15</sup> we are very far from having an adequate account of the psychology of hypothesis formation and confirmation. It is well known that from the purely logical point of view any empirical proposition confirms or disconfirms an infinity of others. For example, (8) is logically equivalent to (9), and any proposition that confirms the latter confirms the former:

(8) All snow is white.

(9) Anything that is not white is not snow.

Proposition (9) is confirmed by anything that is not white and not snow: for example, a green apple. Hence, (10) confirms the claim that all snow is white:

<sup>14</sup> See, for example, Levinson, *Pragmatics*, p. 116; and Brown and Yule, *Discourse Analysis*, pp. 33–5.

<sup>15</sup> J. Fodor, *The Modularity of Mind* (MIT Press, Cambridge, Mass., 1983).

(10) This is a green apple.

Someone who says (10) would not normally, of course, be construed as encouraging the hearer to derive either (8) or (9) as a conclusion. The problem is that there is as yet no principled account capable of explaining which of the infinite set of possible conclusions that the hearer of a given utterance *could* draw will actually be drawn.

Hypothesis formation, according to Fodor, is a creative process involving analogical reasoning, about which virtually nothing is known. Once formed, a given hypothesis will be accepted or rejected on a basis which is again very little understood. As Fodor sees it, the difficulties with hypothesis formation and confirmation arise from the fact that they are *global* as opposed to *local* processes. The distinction between global and local processes corresponds roughly to a distinction between processes which have free access to contextual information and those which do not. A global process is one in which any item of information, however remote and unrelated to the information being processed, may legitimately be used. So, for example, in creating a scientific hypothesis to account for a certain range of data it is legitimate to rely on analogies with other domains of knowledge, seemingly random association of ideas, and any other source of inspiration that comes to hand. Once a hypothesis has been formed, the extent to which it is regarded as confirmed will depend on how well it fits not only with neighbouring domains of knowledge but with one's whole overall conception of the world. A local process, by contrast, is one which needs to take nothing into account apart from the information actually being processed. For example, deductive reasoning from fixed premises is a purely local process in which no attention need be paid to information not contained in the premises themselves. Fodor's argument is that although we have a fair understanding of a variety of local processes, the working of global processes remains a mystery.

Given the haziness of our understanding both of the psychology of hypothesis formation and confirmation and of the effects of context on information processing, it is perhaps not surprising that pragmatists who express scepticism about the role of deductive

background against which this central deductive element is set, and show how 'probabilistic' implicatures would be handled in this framework.

#### RELEVANCE THEORY

Grice's work on the maxims of conversation<sup>18</sup> can be seen as providing the elements of a pragmatic hypothesis confirmation system. Given (a) a source of hypotheses about the speaker's communicative intentions and (b) an adequate account of what it is for a speaker to observe the conversational maxims, it seems reasonable to claim that the most favoured hypothesis, the one the hearer should choose, would be the one which best satisfies the maxims. In at least a few cases, it is easy to see how this proposal might work. Consider disambiguation, for example. Here the source of hypotheses is the grammar, which assigns a range of possible senses to an utterance. On the approach now being considered, the sense the hearer should assume that the speaker wanted to communicate is the one which best accords with the assumption that Grice's maxims have been observed.

Or consider how a hearer might decide which of the deductive consequences of an utterance the speaker wanted to communicate. Here the source of hypotheses is the deductive inference rules. On the approach now being considered, the hearer should assume that the speaker wanted to communicate any subset of consequences needed to satisfy him that the maxims have been observed. Consider, for example, the exchange in (11):

(11) (a) *He*: Is Jacques a good cook?

(b) *She*: He's French, and all the French are good cooks.

Here, in order to satisfy himself that the speaker of (11b) is observing the maxims of relevance and informativeness, the hearer must assume that she wanted him to deduce from her utterance the conclusion in (12):

18 Grice, 'Logic and conversation'.

reasoning in comprehension have said little about the processes by which implicatures are recovered. Bach and Harnish, after establishing that the working out schema is not deductive, add:

Our empirical thinking in general is rife with generalizations and inference principles that we are not conscious of when we use them, if we are conscious of them at all. It would take us well beyond present-day cognitive psychology to speculate on the details of any of this. . . . Whatever these processes are, whatever activates them, whatever principles or strategies are involved, they work, and work well.<sup>16</sup>

But the fact that these processes work well enough in everyday utterance comprehension does not absolve us from saying what they are. If anything, the lack of any existing framework for describing them should make us more, not less, interested in their nature. Given the 'probabilistic' nature of implicatures as illustrated in (2b) above, an adequate theory of how implicatures are recovered might shed light not just on utterance comprehension but on the more general psychological problem of hypothesis formation and confirmation which, for the reasons Fodor has given, has proved to be so intractable.

The claim that pragmatics must be based on *either* deductive *or* non-demonstrative inferential systems, or that the adoption of one type of system must inevitably lead to the rejection of the other, seems to us to be unfounded. For the last few years we have been working on a theory of spontaneous non-demonstrative inference in utterance comprehension and other cognitive domains, in which deductive reasoning plays a central role. In previously published work<sup>17</sup> we have concentrated on the central, deductive element of the theory, which will be outlined briefly in the next section. In the remainder of this paper we will sketch in some of the broader

16 Bach and Harnish, *Linguistic Communication*, p. 83.

17 See, for example, D. Sperber and D. Wilson, 'Mutual knowledge and relevance in theories of comprehension', in *Mutual Knowledge*, ed. N. V. Smith (Academic Press, London, 1982), pp. 61-87; and D. Wilson and D. Sperber, 'On defining relevance', in *Festschrift for Paul Grice*, ed. R. Grandy (OUP, Oxford, forthcoming).

(12) Jacques is a good cook.

In these two types of case, Grice's maxims can thus be used, if not to generate a set of hypotheses about the speaker's communicative intentions, at least to choose among them.

To provide an adequate account of implicature along these lines, two questions would have to be answered. First, what is the source of hypotheses about the possible implicatures of an utterance? Second, is it possible to show more precisely what it is for a speaker to observe Grice's maxims? In previously published work<sup>19</sup> we have offered, in outline, an answer to the second of these questions. We have proposed a definition of relevance and suggested what factors might be involved in assessments of degrees of relevance. We have also argued that all Grice's maxims can be replaced by a single principle of relevance – that the speaker tries to be as relevant as possible in the circumstances – which, when suitably elaborated, can handle the full range of data that Grice's maxims were designed to explain.

We treat relevance as a relation between a proposition *P* and a set of contextual assumptions {*C*}. In previously published work we have made the simplifying assumption that the only propositions used in the comprehension process are those believed to be true. On the basis of this assumption we defined relevance as follows:

(13) A proposition *P* is relevant in a context {*C*} if and only if *P* has at least one contextual implication in {*C*}.

A contextual implication is a special type of logical implication, derived by the use of a restricted set of deductive rules which derive at most a finite set of conclusions from any finite set of premises.<sup>20</sup>

19 For example, D. Wilson and D. Sperber, 'On Grice's theory of conversation', in *Conversation and Discourse*, ed. P. Werth (Croom Helm, London, 1981), pp. 152–77; and Wilson and Sperber, 'On defining relevance'.

20 The details of this system will not concern us here. See, for discussion, D. Sperber and D. Wilson, 'Reply to Gazdar and Good', in Smith, ed., pp. 101–10; Wilson and Sperber, 'On defining relevance'; Sperber and Wilson, *Relevance*.

The contextual implications of a proposition *P* in a context {*C*} are all those conclusions deducible from the union of *P* with {*C*}, but from neither *P* alone nor {*C*} alone. For example, (2b) above contextually implies (3) in a context containing (5), and contextually implies (7) in a context containing (6); it would thus, by our definitions, be relevant in a context containing either (5) or (6).

The intuitive idea behind these definitions is that relevance is achieved when the addition of a proposition to a context modifies the context in a way that goes beyond the mere incrementation of that context with the proposition itself and all its logical implications. As we will show, the production of contextual implications is a special case of a more general notion of contextual modification which emerges once we drop the assumption that the only propositions used in comprehension are those believed to be true. For the moment we will continue to assume that a hearer who wants to establish the relevance of an utterance should be looking for a context with which it will interact to yield contextual implications.

Consider, in this framework, how a hearer might set about processing the information in (14):

(14) *She*: Susan doesn't drink alcohol.

One possible line of interpretation would be to think of the names of some alcoholic drinks, as in (15), and conclude that Susan doesn't drink them, as in (16):

(15) (a) Sherry is alcoholic.  
(b) Gin is alcoholic.  
(c) Whisky is alcoholic.

...

(16) (a) Susan doesn't drink sherry.  
(b) Susan doesn't drink gin.  
(c) Susan doesn't drink whisky.

...



One of the conclusions in (16) might in turn combine with further contextual assumptions to yield a range of further contextual implications, which could in turn combine with further contextual assumptions, and so on indefinitely.

Another line of interpretation would be to think of conditional premises with (14) as antecedent, as in (17), and derive their consequents as conclusions, as in (18):

- (17) (a) If Susan doesn't drink alcohol, she may prefer a soft drink.
- (b) If Susan doesn't drink alcohol, she probably disapproves of getting drunk.
- (c) If Susan doesn't drink alcohol, she never has hangovers.

- (18) (a) Susan may prefer a soft drink.
- (b) Susan probably disapproves of getting drunk.
- (c) Susan never has hangovers.

Again, one of the conclusions in (18) may combine with further contextual assumptions to yield further contextual implications, and so on indefinitely.

We assume that in processing a proposition the hearer begins by systematically searching for contextual implications in a small, immediately accessible context consisting of the propositions that have most recently been processed. To these, further assumptions may be added subject to the following constraint. We assume that information is stored in memory in encyclopaedic entries attached to concepts, and that the information in a given encyclopaedic entry can only be accessed via the presence in the set of propositions currently being processed of the concept to which it is attached. For example, an utterance mentioning alcohol makes accessible (to varying degrees) the set of propositions in the encyclopaedic entry attached to the concept *alcohol*; these in turn give access to the encyclopaedic entries attached to the concepts they contain, and so on indefinitely.

It can be seen that in this framework not all the contextual implications of a given proposition will be equally easy to obtain. Those derived from small, easily accessible contexts will be relatively cheap in processing terms. Those derived from larger, less easily accessible contexts will be relatively expensive in processing terms, because of the additional effort required to access the contexts needed to derive them and to search these contexts systematically for contextual implications. We assume that the universal aim in processing is to obtain the maximum of contextual implications in return for any processing effort expended. However, at a certain point in processing – which will vary from person to person and situation to situation – the cost of obtaining any further contextual implications will become too high, and processing will stop.

Let us say that, other things being equal, the relevance of a proposition increases with the number of contextual implications it yields and decreases as the amount of processing needed to obtain them increases. Maximizing the relevance of a proposition is thus a matter of accessing, as quickly as possible, a context in which it will yield the maximum of contextual implications in return for the available processing effort. The most relevant propositions will be those which yield a wide range of contextual implications in a small, immediately accessible context.

We assume that the universal goal in cognition is to acquire relevant information, and the more relevant the better. We also assume that a speaker who thinks it worth speaking at all will try to make his utterance as relevant as possible. A hearer should therefore bring to the processing of every utterance the standing assumption that the speaker has tried to be as relevant as possible in the circumstances. It is this assumption that we call the principle of relevance.

A speaker cannot observe the principle of relevance without believing that his utterance will convey some relevant information to the hearer. Sometimes, he may have only the most general grounds for thinking so. For example, if I know you follow the pop music charts, I can reasonably assume that it will be relevant to you to know the name of the new number one hit, even though

I may have no idea what specific implications this information will have for you. At other times, however, a speaker may have a much more specific idea of the sort of context that will be brought to bear and the sort of conclusions derived. It is in situations like this that we believe implicatures arise.

Consider, for example, the exchange in (19):

- (19) (a) *He*: Does Susan drink whisky?  
 (b) *She*: She doesn't drink alcohol.

On what grounds might the speaker of (19b) have thought her utterance would be relevant to the hearer? What sort of context might she have expected him to supply that would be both accessible enough and rich enough in contextual implications for it to be worth his while to process her utterance? The answer is clear. He has just asked her whether Susan drinks whisky. In our framework, he would not have asked this question if he had not had immediately accessible a context in which the information that she did (or did not) drink whisky would be relevant – and indeed more relevant than any other information he thinks she will be able to provide. By providing this information directly, she would therefore be sure of satisfying the principle of relevance.

In fact, her utterance does not provide the information directly: the hearer has first to supply the contextual assumption in (15c) and then to derive the conclusion in (16c):

- (15) (c) Whisky is alcoholic.  
 (16) (c) Susan doesn't drink whisky.

However, the speaker can reasonably expect him to do this. On the one hand, her utterance gives him immediate access to his encyclopaedic entry for *alcohol*, which should in turn provide access to propositions of the form (15a–c). On the other hand, on normal assumptions about the organization of memory, the immediately preceding mention of whisky should act as a prompt, making (15c) more accessible than other propositions of this form. It would therefore be reasonable to assume that one of the grounds

on which the speaker of (19b) thought her utterance would be relevant was that she expected it to be processed in a context which contained (15c) as an assumption and yielded (16c) as a contextual implication.

We want to say that the speaker of (19b) implicates both (15c) and (16c). On this approach, the implicatures of an utterance are those contextual assumptions and implications which the hearer has to recover in order to satisfy himself that the speaker has observed the principle of relevance. Here, (15c) is a necessary precondition on the recovery of (16c), and (16c) is a necessary precondition on the recovery of the whole range of contextual implications on which the main relevance of (19b) depends. We will call (15c) an *implicated assumption* and (16c) an *implicated conclusion* of (19b).

In fact, as we have described it, the interpretation of (19b) does not conform to the principle of relevance. The speaker could have conveyed the whole of this interpretation more economically by producing the direct answer (16c). Instead, she has forced the hearer to process the proposition expressed by (19b), to access (15c) and to deduce (16c) as a contextual implication, each step requiring some processing effort which would not have been required by the direct answer (16c). Suppose the hearer asks himself why she might have thought that the indirect answer (19b) would be more relevant to him than a direct answer. The only possible explanation is that she must have expected it to yield some additional contextual implications, not derivable from the direct answer (16c), which would more than compensate for the extra processing cost. In other words, the only possible explanation is that she believed that the surplus of information she was providing had some relevance in its own right.

As always, the speaker must have some reason for thinking that this surplus of information will be relevant, and more relevant than any alternative information she could provide. She may know, for example, that the hearer is wondering what drink to offer Susan. In these circumstances, her response in (19b) would encourage him to derive conclusions along the lines of (16) and (18c) above: that she doesn't drink sherry, gin, etc., and that she may prefer a soft drink. Or he may be wondering whether to invite Susan to his

party. In these circumstances, the response in (19b) would encourage him to derive conclusions along the lines of (18b) above: that Susan may disapprove of getting drunk, that he should maybe not bother inviting her to his party, and so on.

Grice suggests that many implicatures are indeterminate:

Since to calculate a conversational implicature is to calculate what has to be supposed in order to preserve the supposition that the Co-Operative Principle is being observed, and since there may be various possible specific explanations, a list of which may be open, the conversational implicature in such cases will be a disjunction of such specific explanations; and if the list of these is open, the implicature will have just the kind of indeterminacy that many actual implicata do in fact seem to possess.<sup>21</sup>

His commentators have been divided in their reaction to this suggestion. Some, realizing the difficulty of providing an explicit treatment of indeterminacy, have largely ignored it. Gazdar, for example, notes the existence of indeterminacy, but adds: 'Because indeterminacy is hard to handle formally, I shall mostly ignore it in the discussion that follows. A fuller treatment of implicatures would not be guilty of this omission, which is really only defensible on formal grounds.'<sup>22</sup> Others, less interested in explicit treatment of the processes by which implicatures arise, tend to use the indeterminacy of implicatures as an argument against deductive models of the recovery process and in favour of 'informal', 'loose' or 'probabilistic' models.<sup>23</sup> In our framework, the indeterminacy of implicatures can be dealt with without losing the explicitness of the deductive approach.

Sometimes, as we have shown, a speaker can observe the principle of relevance without having any idea of the sort of context the utterance will be processed in, or the sort of conclusions that will be derived. In these cases, the utterance will have no implicatures at all. In other cases, as with (19b) and its implicatures

<sup>21</sup> Grice, 'Logic and conversation', p. 58.

<sup>22</sup> G. Gazdar, *Pragmatics: Implicature, Presupposition and Logical Form* (Academic Press, New York, 1979).

<sup>23</sup> See, for example, Leech, *Principles of Pragmatics*, chapters 2 and 7.

(15c) and (16c), it is impossible to see how the speaker could have observed the principle of relevance without expecting a specific contextual assumption to be supplied and a specific conclusion derived. In these cases the utterance will have fully determinate implicatures. Between these two extremes lie a whole range of intermediate cases. We have discussed a situation where the indirect answer (19b) would encourage the hearer to think of assumptions along the lines of (15) or (17) and conclusions along the lines of (16) or (18). Here, the speaker has a general idea of the type of assumption to be supplied and the type of conclusion to be derived, but may not know or care which specific assumptions and conclusions of this type will be supplied. The clearer an idea the speaker must have had of the specific assumptions and conclusions to be supplied, the more determinate the implicatures will be; the vaguer an idea he could have had and still have been observing the principle of relevance, the less determinate the implicatures will be, up to the point where they vanish altogether and the choice of contexts and conclusions is left solely up to the hearer.

In every case, the method of processing is the same. The hearer supplies specific contextual assumptions and derives specific contextual implications. What varies is not the specificity of the assumptions and conclusions derived, or the formality of the reasoning processes involved, but simply the amount of foreknowledge the speaker must be taken to have had of the way the utterance would be processed, and with it the degree of responsibility he must take for the particular conclusions derived. Suppose, for instance, that I could not have observed the principle of relevance without expecting you to supply a certain assumption and derive a certain conclusion in processing my utterance. Then by encouraging you to supply them, I take as much responsibility for their truth as for the truth of the proposition I have explicitly expressed.

To take an example, suppose that after the exchange in (19) has taken place it turns out that Susan does drink whisky. Although (19b) does not *entail* that Susan does not drink whisky (the additional assumption that whisky is alcoholic is needed), the speaker could be quite rightly accused of having misled the hearer by allowing him to suppose that she did. Similarly, a speaker who

secretly believed that whisky was not alcoholic could be accused at least of having *tried* to mislead the hearer by uttering (19b) and thus encouraging him to suppose that whisky was alcoholic. In other words, the speaker is committed to the truth of all determinate implicatures conveyed by her utterance, just as much as if she had expressed them directly.

With less determinate implicatures, the speaker cannot be held solely responsible for their truth. Suppose, for example, that the exchange in (19) takes place in the circumstances described above, where the speaker of (19a) is wondering whether to invite Susan to his party. Here, (19b) would clearly carry implicatures to the effect that Susan may disapprove of getting drunk, is unlikely to enjoy rowdy behaviour, may inhibit the proceedings, and is thus not a suitable person to invite to a party. However, it would be a little strong to say that the speaker of (19b) had specifically indicated that the hearer should not invite Susan to his party: this is only one among a range of roughly equivalent conclusions that the hearer could have drawn, any of which would have satisfied him that the speaker had observed the principle of relevance by providing some information with a bearing on his deliberations. The weaker the implicature – that is, the wider the range of roughly equivalent alternative assumptions and conclusions that would have satisfied the hearer that the speaker had observed the principle of relevance – the weaker the speaker's responsibility for its truth, up to the point where the implicature disappears altogether and the responsibility for the assumptions used and the conclusions drawn from them lies solely on the side of the hearer.

Talk of degrees of responsibility for the truth of implicated assumptions and conclusions takes us outside the simplified framework we have been assuming so far – a framework which abstracts away from the fact that a proposition can be expressed by the speaker with a stronger or weaker guarantee of truth, and that this guarantee may be more or less trusted by the hearer.<sup>24</sup> In the next section we will outline only as much as is needed to account for the 'probabilistic' nature of implicatures.

<sup>24</sup> The full framework is developed in detail elsewhere. See Sperber and Wilson, *Relevance*.

### AN EXTENSION TO THE THEORY

What would happen, in our simplified framework, if the hearer of (20) tried to process it in a context such as (21a–d), which directly contradicts it?

(20) *Sbe*: Peter is not coming to the party.

(21)(a) Peter is coming to the party.

(b) If Peter is coming to the party, Jane will come.

(c) Jane will come to the party.

(d) If Peter is not coming to the party, Harry will not come.

We assume that no contextual implications at all are derivable from a contradictory set of assumptions.<sup>25</sup> The hearer of (20) must therefore either reject the utterance as irrelevant, or modify his assumptions in (21). By eliminating (21a) he could, in the simplified system outlined above, establish the relevance of (20) in a context consisting of the remaining assumptions (21b–d), deriving (22) as a contextual implication:

(22) Harry will not come to the party.

However, this is a rather unsatisfactory account, for two reasons. First, it implies that if (21d) had not been present in the context to enable (22) to be derived, (20) could not have been relevant at all. Yet intuitively it is always relevant to discover that one has been mistaken. In such cases, our original intuition that a proposition achieves relevance by interacting with, or modifying, the context in which it is processed, is not matched by our formal definition. Second, it is clear that a hearer in real life might neither reject (20) and retain (21a) nor accept (20) and reject (21a). He might decide that on the whole (20) is more likely to be true than (21a), or that (21a) is more likely to be true than (20), or that he has no idea which of (20) and (21a) is true. To account for these facts we must abandon the assumption that the only way a proposition can

<sup>25</sup> See *ibid.* for discussion.

modify a context is by yielding contextual implications, and that the only propositions which play a role in processing are those regarded as certainly true.

A proposition may be put forward, and accepted, with varying degrees of confidence. Some account of how this can happen must be provided by any adequate theory of cognition. We argue elsewhere that a logical theory of confirmation, involving the assignment of numerical confirmation values, is not the best starting point for such an account.<sup>26</sup> For the purposes of this paper, however, let us say that when a proposition is processed it is assigned some form of gross absolute confirmation value representing its estimated likelihood of being true. Positive values represent the estimate that it is more likely to be true than false; the highest positive value, *true*, represents it as certainly true. Negative values represent the estimate that it is more likely to be false than true; the lowest negative value, *false*, represents it as certainly false. The absence of a value represents the absence of an opinion either way. When (20) is added to the context in (21), some complementary assignment of values to (20) and (21a) must be achieved: for example, by giving (20) the value *true* and (21a) the value *false*, or by giving (20) some positive value less than *true* and (21a) some negative value greater than *false*. The greatest possible effect that (20) could have would be to make the hearer entirely abandon his former assumption (21a), assigning (20) the value *true* and (21a) the value *false*. Let us assume that this readjustment takes place.

The deductive rules must now be applied to a context in which one premise, (21a), is *false*. Moreover, (21a) combines with (21b) to yield (21c) as a conclusion. What happens to this conclusion now that one of the premises used to derive it is *false*? Let us say that only conclusions based on premises with positive confirmation values can have confirmation values of their own, so that once (21a) is assigned the value *false*, (21c) automatically loses its own value. Any further conclusions involving (21c) will in turn be affected. Let us also assume that (21d) has some positive value less than *true*. How does this affect the value of (22), a contextual

<sup>26</sup> Ibid.

implication based on (20) and (21d) as premises? Let us say (simplifying slightly) that a conclusion based on a mixture of premises with positive values will inherit at most the lowest value of any premise used in deriving it. Thus (22) will inherit at most the value of (21d).

The total effect of a proposition on a context can now be assessed by answering the following questions. First, did it directly affect the value of any proposition already present in the context, as (20) affected the value of (21a)? If so, how large was the change? Second, did it indirectly affect the value of any proposition already present in the context, as the modification of (21a) indirectly affected the value of (21c)? If so, how many propositions were affected, and how large was the change? Third, did its addition to the context yield any new contextual implications, as (20) yielded (22)? If so, how many were there, and how high were their values? We assume that the higher the value of any new contextual implication, the greater the modification to the context, and that a new contextual implication which lacks a confirmation value does not modify the context at all. With this extension to the framework, the recovery of contextual implications becomes just a special case of a more general notion of contextual modification in terms of which relevance can be redefined.

Let us say, then, that the proposition *P* is relevant in a context {*C*} if and only if it modifies {*C*} in one of the ways described above. Let us say that other things being equal, the more *P* modifies the context the more relevant it is, but that other things being equal, the greater the amount of processing required to bring about this modification, the less relevant it is. As before, the aim of the hearer in processing *P* will be to access a context which makes the best possible use of the available processing resources – that is, a context that maximizes the relevance of *P*.

From this very brief account, one or two general principles emerge. In particular, the only new contextual implications worth deriving will be those based on assumptions with positive values, the higher the better. The use of other assumptions will incur processing costs, but without leading to any reward in terms of contextual modification, since no new contextual implication based on them will be assigned a confirmation value at all. The use

of such assumptions will thus detract from the relevance of any proposition being processed. Retrieval strategies geared to the maximization of relevance should therefore be aimed at retrieving, as quickly as possible, assumptions with positive confirmation values, the higher the better, and this fact should be known to any speaker. The hearer should also be able to infer, from the fact that the speaker is observing the principle of relevance, that he must have believed that all implicated conclusions, and all assumptions needed to derive them, had positive confirmation values, even if the hearer, of his own knowledge, would have been inclined to treat them as *false*.

Let us return, in the light of this discussion, to (2b) and its implicature (3):

- (2) (a) *He*: Will you have some coffee?  
 (b) *She*: Coffee would keep me awake.  
 (3) She won't have any coffee.

We can now answer the two questions raised earlier: how is (3) derived, and how is the alternative derivation of (7) ruled out?

- (7) She will have some coffee.

The answer to the first question is that (3) is a contextual implication of (2b) in a context containing (5):

- (5) (a) She doesn't want to be kept awake.  
 (b) She won't have anything that would keep her awake.

Moreover, (5a-b) are implicated assumptions, and (3) is an implicated conclusion. By parallel arguments to those used in the discussion of relevance theory, the hearer knows that a speaker observing the principle of relevance must have expected him to supply the assumptions in (5) and derive the conclusion in (3). A question we have not yet considered is what confirmation value the speaker expects the hearer to assign to the implicated assumptions and conclusion.

If the exchange was taking place in the evening, the hearer could no doubt, of his own knowledge, assign some positive confirmation value less than *true* to (3) and (5). Under normal assumptions, it would be no more than a probability that the speaker does not want to be kept awake, and will not have anything that would keep her awake, and hence no more than a probability that she won't have any coffee. In the circumstances, however, the information that the speaker probably won't have any coffee would not be relevant enough. In the first place, she ought to know whether she will or will not have any coffee, and if she knows, she should have told him, since he has indicated that this information would be relevant to him. In the second place, she must realize that the hearer would have been aware that she *might* not want any coffee when he asked his question. What he indicated by asking it was that a categorical answer would be relevant to him. Therefore, if she has this information and is observing the principle of relevance, she ought to have given it. Is there any way of construing (2b) as giving a categorical answer to the question in (2a)? Certainly: all the hearer has to do is upgrade the values of (3) and (5) to *true*. To preserve the assumption that the speaker has observed the principle of relevance, this is what he must do. By the arguments of relevance theory, the speaker, who has encouraged him to do this, will be held just as responsible for the truth of (3) and (5) as if she had expressed these propositions directly.

To complete the interpretation, some justification has to be found for the fact that the speaker has chosen an indirect rather than a direct form of answer. Failure to find such a justification would be *prima facie* evidence that this line of interpretation was not correct. Here, at least on an informal level, the reason is easy to see. A direct refusal, with no explanation, would be likely to raise all sorts of questions in the hearer's mind about why his offer has been refused. The indirect answer (2b) simultaneously refuses the offer of coffee and explains the refusal, thus saving the hearer the time he might have spent speculating on the reasons behind it. This line of interpretation is thus confirmed as satisfying the principle of relevance.

By contrast, an interpretation based on the contextual assumptions in (6) and the conclusion in (7) is unlikely to be considered at

all. If considered, it should be rejected as not conforming to the principle of relevance:

- (6) (a) She wants to be kept awake.
- (b) She will have anything on offer that would keep her awake.
- (7) She will have some coffee.

In the first place, if the exchange was taking place in the evening, the hearer of (2b) would no doubt, of his own knowledge, assign a lower confirmation value to (6) than to (5). On the assumption that retrieval strategies give preferential access to higher-valued assumptions, (5) should be retrieved before (6); and since it gives rise to a satisfactory interpretation, there is no reason why (6) should not be considered at all. However, suppose it is. The interpretation it gives rise to should still be rejected as not conforming to the principle of relevance. There is no reason why a speaker observing the principle of relevance should have preferred the indirect answer (2b), construed in this way, to the direct answer that she wants some more coffee. An acceptance, unlike a refusal, normally needs no justification; it normally raises no questions in a hearer's mind about the reasons behind it. Moreover, a speaker who attempted to justify her acceptance on the grounds that she wanted to stay awake and would accept anything on offer that would keep her awake *would* raise a number of questions in her hearer's mind, and cost him valuable processing time if she did not go on immediately to answer them. Hence, the line of interpretation based on (6) and (7) is ruled out at a number of points as failing to conform to the principle of relevance.

Many standard examples of implicature fit quite straightforwardly into this framework. For example, Clark discusses a class of 'bridging' implicatures needed to establish the reference of the referring expressions in (23b)–(25b):<sup>27</sup>

- (23) *She*: (a) I went into the room. (b) The window was open.
- (24) *She*: (a) I went into the room. (b) Both windows were open.
- (25) *She*: (a) I went into the room. (b) All three windows were open.

27 H. H. Clark, 'Bridging', in *Thinking: Readings in Cognitive Science*, eds P. N. Johnson-Laird and P. C. Wason (CUP, Cambridge, 1977), pp. 411–20.

As Clark points out, in normal circumstances the hearer of these utterances would supply assumptions (26)–(28) respectively, even if there had been no previous mention of the number of windows in the room:

- (26) The room had a window.
- (27) The room had two windows.
- (28) The room had three windows.

Consider how this might happen in the case of (24b). Given the immediately preceding mention of a room, the hearer of (24b) would no doubt, of his own knowledge, have relatively easy access to each of the assumptions in (26)–(28), and be able to assign each of them some confirmation value less than *true*. On the assumption that the speaker has observed the principle of relevance, he will take it that he must upgrade the value of (27) to *true* and use it to establish the reference of the referring expression *both windows* in (24b). As long as this assignment leads to a satisfactory range of contextual implications, he will accept it as correct. The role of the implicated assumption here is not to yield any particular contextual implication, but to establish the referential content of the utterance, which is a necessary precondition to recovering any contextual implications at all.

There may, of course, be many other logically possible assumptions that the hearer of (24b) could have used: for example, the windows might have been in the house opposite, or mentioned in a letter the speaker found in the room. However, in normal circumstances, unless the existence of these windows had already been established, these assumptions would be much less accessible than those in (26)–(28) above, and a speaker observing the principle of relevance could not normally have expected the hearer to supply them. The general principle, for bridging implicatures as for all other implicated assumptions, is that they should – at least in the estimation of the speaker – be virtually instantaneously accessible, and more accessible than any alternative assumption likely to lead to an acceptable interpretation. If not, a speaker observing the principle of relevance should have done something to increase their accessibility – for example by directly mentioning

them in the utterance – and thus save the hearer some unnecessary processing costs.

Grice is unsure whether the implicatures carried by (29b) would be categorical or merely probable:

- (29)(a) A: I am out of petrol.  
 (b) B: There is a garage round the corner.

He claims that 'B would be infringing the maxim 'Be relevant' unless he thinks, or thinks it possible, that the garage is open and has petrol to sell; so he implicates that the garage is, or at least may be open, etc.<sup>28</sup>

We can shed some light on his uncertainty. A speaker observing the principle of relevance should expect the hearer, among other things, to supply the contextual assumption in (30) and derive the conclusion in (31):

- (30) If there's a garage round the corner, I can get some petrol there.  
 (31) I can get some petrol round the corner.

However, if it occurs to the hearer that garages may be closed or out of petrol, he will be unable of his own knowledge to assign more than a fairly high degree of probability to (30) or (31). In other words, he will only be able to derive the conclusion that he *may* be able to get some petrol round the corner. So far, the case is exactly like the two previous ones discussed in this section. The difference is that in this case, even the information that he *may* be able to get petrol round the corner might be relevant enough for a speaker who could not make a more categorical claim to think this information worth offering.

More precisely, what the speaker of (29b) indicates is that *as far as he knows* (30) and (31) are true, and therefore that *as far as he knows* the garage is open and selling petrol, etc. In some circumstances – for example if the speaker is coming from the direction of the garage with a full petrol can in his hand – the hearer would be

<sup>28</sup> Grice, 'Logic and conversation', p. 51.

justified in assuming that he knew with certainty; in this case the implicatures (30) and (31), and their necessary conditions that the garage is open and selling petrol, would be regarded as categorical. In different circumstances only a weaker attribution of confirmation value would be justified. The general principle is thus that when a less than categorical implicature would conform to the principle of relevance, the hearer is not entitled to assume that the speaker expected a categorical implicature to be supplied.

All the implicatures considered so far have arisen in processing the explicit content of the utterance, or in Grice's terms 'what was strictly speaking said'. As Grice points out, many implicatures arise not so much from the content of what was said as from the saying of it, in those circumstances, to that audience, and so on. Consider (32), for example:

- (32) (a) A: Where does C live?  
 (b) B: Somewhere in the South of France.

In normal circumstances, B would implicate that he does not know more precisely where C lives. We would analyse this implicature as resulting not from the explicit content of (32b) but from the fact, which the hearer is expected to notice and process, that the speaker has failed to give more precise information when, in the circumstances, more precise information would have been more relevant. To reconcile this fact with the assumption that the speaker has observed the principle of relevance, the hearer would have to access something like the contextual assumption in (33) and derive something like the conclusion in (34):

- (33) If B has failed to give more precise information on C's whereabouts, he has no more precise information to give.  
 (34) B has no more precise information to give.

Here, (34) is a contextual implication, not of the content of (32b), but of (35):

- (35) B has failed to give more precise information on C's whereabouts.



This is a proposition which may, and in this case clearly does, have some relevance in its own right. It would therefore be a mistake to think of a speaker as providing relevant information only through the explicit content of his utterance. An act of utterance may draw the hearer's attention to a number of propositions other than explicitly expressed, and these may contribute to the overall relevance of the utterance.

Grice analyses (32b) as involving a clash between the maxims of informativeness and truthfulness; the desire to give precise information about *C*'s whereabouts is sacrificed to the demands of truthfulness. He says little to explain why the supposed clash is not resolved in the opposite direction, which is a weakness of his system. Any system with more than one pragmatic principle must provide some account of their interaction – an account which is rarely provided. In our framework, with its single principle, there is no possibility of clashes. Notice, too, that there is no appeal to a violation of the principle of relevance, real or apparent, in our account of (32b). The speaker has been as relevant as he can in the circumstances – more relevant, for example, than if he had merely said 'I don't know.' His failure to provide more detailed information is explained not by appeal to deliberate violation of the principle of relevance, but by the assumption that he did not have more detailed information to provide.

We would like to make the more general claim that the recovery of implicatures *never* involves an appeal to deliberate violation of the principle of relevance. Grice himself remarks that it is hard to find cases in which his maxim of relevance is deliberately violated, but offers the following as a candidate example:

At a genteel tea party, *A* says 'Mrs *X* is an old bag.' There is a moment of appalled silence, and then *B* says 'The weather has been quite delightful this summer, hasn't it?' *B* has blatantly refused to make what *he* says relevant to *A*'s preceding remark. He therefore implicates that *A*'s remark should not be discussed and, perhaps more specifically, that *A* has committed a social gaffe.<sup>29</sup>

However, the fact that *B*'s utterance is not relevant to the immediately preceding remark does not mean that it is not relevant

<sup>29</sup> *Ibid.*, p. 54.

at all. Most of its relevance would be achieved, not through its content, but by drawing the hearers' attention to the fact that *B* is deliberately ignoring *A*'s remark. To reconcile this with the assumption that the principle of relevance has been observed, they would have to access assumptions along the lines of (36a–b) and derive conclusions along the lines of (37a–b):

- (36) (a) If *B* is deliberately ignoring *A*'s remark, he believes it should not be discussed.  
 (b) If *B* believes *A*'s remark should not be discussed, he believes it was a social gaffe.
- (37) (a) *B* believes *A*'s remark should not be discussed.  
 (b) *B* believes *A*'s remark was a social gaffe.

In our framework these would be implicatures, since they are needed to reconcile the fact that *B* is deliberately ignoring *A*'s remark with the assumption that he is being as relevant as he can. They would, however, be relatively weak implicatures, since a variety of roughly equivalent assumptions could be made, all of which would reconcile *B*'s behaviour with the assumption that he is observing the principle of relevance. By the arguments outlined above, the speaker could therefore not be held solely responsible for their truth.

Relevance theory thus makes a number of specific claims about the role of implicatures in comprehension and the processes by which they are recovered. First, they are either contextual assumptions or contextual implicatures which the hearer is expected to supply in satisfying himself that the speaker has observed the principle of relevance. Implicated assumptions are recovered by the same processes used to retrieve other contextual assumptions, with ease of accessibility playing a decisive role; implicated conclusions are recovered by deduction.

Second, the hearer may be expected to upgrade the confirmation value of an implicated assumption to the point where the conclusion it yields conform to the principle of relevance. The speaker is held responsible for the truth (or degree of confirmation) of any assumptions and conclusions upgraded in this way. The hearer may thus acquire new information, not only from the explicit

This is a proposition which may, and in this case clearly does, have some relevance in its own right. It would therefore be a mistake to think of a speaker as providing relevant information only through the explicit content of his utterance. An act of utterance may draw the hearer's attention to a number of propositions other than explicitly expressed, and these may contribute to the overall relevance of the utterance.

Grice analyses (32b) as involving a clash between the maxims of informativeness and truthfulness; the desire to give precise information about *C*'s whereabouts is sacrificed to the demands of truthfulness. He says little to explain why the supposed clash is not resolved in the opposite direction, which is a weakness of his system. Any system with more than one pragmatic principle must provide some account of their interaction – an account which is rarely provided. In our framework, with its single principle, there is no possibility of clashes. Notice, too, that there is no appeal to a violation of the principle of relevance, real or apparent, in our account of (32b). The speaker has been as relevant as he can in the circumstances – more relevant, for example, than if he had merely said 'I don't know.' His failure to provide more detailed information is explained not by appeal to deliberate violation of the principle of relevance, but by the assumption that he did not have more detailed information to provide.

We would like to make the more general claim that the recovery of implicatures *never* involves an appeal to deliberate violation of the principle of relevance. Grice himself remarks that it is hard to find cases in which his maxim of relevance is deliberately violated, but offers the following as a candidate example:

At a genteel tea party, *A* says 'Mrs *X* is an old bag.' There is a moment of appalled silence, and then *B* says 'The weather has been quite delightful this summer, hasn't it?' *B* has blatantly refused to make what *he* says relevant to *A*'s preceding remark. He therefore implicates that *A*'s remark should not be discussed and, perhaps more specifically, that *A* has committed a social gaffe.<sup>29</sup>

However, the fact that *B*'s utterance is not relevant to the immediately preceding remark does not mean that it is not relevant

<sup>29</sup> *Ibid.*, p. 54.

at all. Most of its relevance would be achieved, not through its content, but by drawing the hearers' attention to the fact that *B* is deliberately ignoring *A*'s remark. To reconcile this with the assumption that the principle of relevance has been observed, they would have to access assumptions along the lines of (36a–b) and derive conclusions along the lines of (37a–b):

- (36) (a) If *B* is deliberately ignoring *A*'s remark, he believes it should not be discussed.  
 (b) If *B* believes *A*'s remark should not be discussed, he believes it was a social gaffe.  
 (37) (a) *B* believes *A*'s remark should not be discussed.  
 (b) *B* believes *A*'s remark was a social gaffe.

In our framework these would be implicatures, since they are needed to reconcile the fact that *B* is deliberately ignoring *A*'s remark with the assumption that he is being as relevant as he can. They would, however, be relatively weak implicatures, since a variety of roughly equivalent assumptions could be made, all of which would reconcile *B*'s behaviour with the assumption that he is observing the principle of relevance. By the arguments outlined above, the speaker could therefore not be held solely responsible for their truth.

Relevance theory thus makes a number of specific claims about the role of implicatures in comprehension and the processes by which they are recovered. First, they are either contextual assumptions or contextual implications which the hearer is expected to supply in satisfying himself that the speaker has observed the principle of relevance. Implicated assumptions are recovered by the same processes used to retrieve other contextual assumptions, with ease of accessibility playing a decisive role; implicated conclusions are recovered by deduction.

Second, the hearer may be expected to upgrade the confirmation value of an implicated assumption to the point where the conclusions it yields conform to the principle of relevance. The speaker is held responsible for the truth (or degree of confirmation) of any assumptions and conclusions upgraded in this way. The hearer may thus acquire new information, not only from the explicit

content of an utterance and its implicated conclusions, but also from its implicated assumptions.

Third, implicatures may be recovered either in the course of processing the explicit content of an utterance, or in the course of processing some higher-level description of it which the hearer is expected to construct. It is important, therefore, that relevance is defined not as a relation between utterances, but as a relation between propositions or sets of propositions.

Fourth, there is no essential connection between the recovery of an implicature and the assumption that the speaker has deliberately violated the principle of relevance. On the contrary: in this aspect of comprehension as in every other, it is the assumption that the speaker has *not* violated the principle of relevance that makes all the difference between processing an item of information that has been deliberately communicated and one that has not.

Within this framework, utterance comprehension is ultimately a matter of hypothesis formation and confirmation: the best hypothesis about the speaker's communicative intentions is the one that best satisfies the principle of relevance. However, it does not follow that deductive inference plays no role in the formation and confirmation of pragmatic hypotheses. On the contrary, because relevance is itself defined in partly deductive terms, the description of pragmatic hypothesis formation and confirmation makes essential reference to deductive processing. In particular, the class of possible implicated assumptions must be of a form capable of combining with information derived from the utterance to undergo deductive inference rules, and the class of possible implicated conclusions is itself deductively determined. The assumption that the overall framework in which comprehension takes place is ultimately non-demonstrative is not incompatible with the assumption that deductive processing plays a central role in comprehension.

#### CONCLUDING REMARK

As we have described them, the processes of pragmatic hypothesis formation and confirmation are clearly context dependent. But to

what extent are they global processes in Fodor's sense? They are global in principle because, as we have shown, the interpretation of a given utterance can proceed, in ever expanding contexts, just as long as the hearer thinks that the rewards are likely to outweigh the processing costs. With certain types of utterance, for example a sacred text or a fortune-teller's prophecy, a hearer might be willing to devote a lifetime's effort to the interpretation process. In practice, however, expectations of relevance are generally much lower; there are other demands on the hearer's processing time; and he generally satisfies himself with establishing relevance in the most immediately accessible context, and leaves it at that.

The same point applies to the formation and confirmation of scientific hypotheses. At one extreme are major theories, which may take a lifetime to develop and confirm. At the other extreme are such minor hypotheses as that there is a bird on the grass, or that spring is on the way, which are formulated in passing, processed in the most immediately accessible context and either abandoned or stored for future use.

We see little reason to think that there are differences of principle between relatively local and relatively global processes of pragmatic or scientific hypothesis formation and confirmation. In particular, the role of context and the goal of maximizing relevance are the same for both types of process, although the principle of relevance, of course, applies only to deliberately communicated information. What distinguishes the two types of process is simply a relative difference, in practice not in principle, in their freedom of access to contextual information. We therefore suggest that the most useful way of gaining insight into fully global processes in Fodor's sense is to look at the relatively more local processes of everyday utterance interpretation and the interpretation of everyday sights and sounds, and work on the assumption that other, more global processes of hypothesis formation and confirmation involve the same cognitive principles and goals.