

analysis, the speaker of (157) and (158) is guaranteeing that Jane is cleverer than the hearer would otherwise have expected. Thus, the intuition that exclamatives are emphatic assertions and the striking parallelisms between exclamative and interrogative form are simultaneously explained.

This very sketchy discussion of speech acts illustrates the general relevance of the principle of relevance. The principle of relevance makes it possible to derive rich and precise non-demonstrative inferences about the communicator's informative intention. With the principle, all that is required is that the properties of the ostensive stimulus should set the inferential process on the right track; to do this they need not represent or encode the communicator's informative intention in any great detail. Thus, illocutionary-force indicators such as declarative or imperative mood or interrogative word order merely have to make manifest a rather abstract property of the speaker's informative intention: the direction in which the relevance of the utterance is to be sought.

Postface

1 Introduction

In the nine years since *Relevance* was first published, the theory of communication it proposes has been widely accepted, widely criticised and widely misunderstood. The book has been translated into several languages;¹ its implications for pragmatic theory have been explored in a growing number of books and articles; it has inspired work in neighbouring disciplines, including linguistics, literary studies, psychology and philosophy. In section 2 of this postface, we review briefly the main developments that have taken place since the first edition was published.²

Many commentators, to whom we are very grateful, have raised a wide variety of objections to the theory.³ We have had the opportunity to answer most of them in a series of publications to which interested readers are referred.⁴ These criticisms have helped us correct some mistakes in the book; they have also made us aware of the difficulties in comprehension and the many possibilities of misunderstanding it presents. Either because we are dense, or because we have had more time than our commentators to think about these issues, we find that the most serious problems with our theory are those we have discovered ourselves. In section 3 of this Postface, we outline these problems, and propose several significant changes both of formulation and of substance.

2 Developments

There is now a substantial body of work expounding and evaluating the basic ideas of relevance theory. This includes a précis of *Relevance*,⁵ two textbooks and large sections of an encyclopaedia of pragmatics,⁶ expository articles designed for non-specialist audiences,⁷ and several lengthy

critiques.⁸ The implications of the theory have been explored in monographs and dissertations,⁹ there are edited collections including papers in relevance theory,¹⁰ there is an e-mail network for exchange of ideas, and a bibliography for classroom use.¹¹ Several research projects have been undertaken; informal workshops are held each year in London, and more formal conferences and lecture series have been held around the world. We will not attempt here a survey of this very diverse literature, but merely point out some of the directions in which we feel that particularly interesting and fruitful work is being done.

2.1 *Explicit communication and the explicit-implicit distinction*

Grice seems not to have noticed (or at least not to have developed the idea) that his Co-operative Principle and maxims could help with other aspects of pragmatic interpretation than the recovery of implicatures: with disambiguation and reference assignment, for example, which he saw as contributing not to what is implicated but to what is (explicitly) said. In 'Logic and Conversation', he gives the impression that sentence meaning and contextual factors are enough on their own to account for disambiguation and reference assignment, and most Gricean pragmatists simply followed him on this.¹² This oversight had two important consequences. First, Gricean pragmatists were slow to react to the extensive psycholinguistic work being done on disambiguation and reference assignment.¹³ Second, they tended to take for granted that pragmatic principles make no contribution to explicit content, and that any aspect of utterance interpretation in which pragmatic principles play a role is automatically an implicature.¹⁴

In *Relevance* (chapter 4, section 2), we rejected this view of pragmatics as *de facto* co-extensive with the study of implicatures. We introduced a notion of explicature, parallel to Grice's notion of implicature, and a definition of explicit communication, which we saw as 'richer, more inferential and hence more worthy of pragmatic investigation than do most pragmatists in the Gricean tradition'. A start was made on studying disambiguation and reference assignment from a relevance-theoretic perspective. We also questioned Grice's suggestion (1989: 25) that disambiguation and reference assignment are the only context-dependent processes involved in explicit communication, drawing attention to a range of further inferential processes required to complete the interpretation of semantically incomplete expressions, narrow the interpretation of vague expressions and, more generally, enrich the linguistically encoded meaning to a point where the resulting overall interpretation would be relevant enough.

The distinction between explicit and implicit communication, and the

role of pragmatic factors in explicit communication, have been the focus of much recent research. As noted in *Relevance* (chapter 4, section 3), psycholinguists have provided valuable insights into the actual processes of disambiguation and reference assignment by investigating, for example, how many candidate interpretations are activated, and at what point one is selected and the others dismissed. However, they have been less interested in what makes the selected interpretation pragmatically acceptable, and on this point pragmatic theorists have a contribution to make. Relevance theory claims that in disambiguation and reference assignment, as in every other aspect of interpretation, the first interpretation consistent with the principle of relevance is the one the hearer should choose.¹⁵ This is not the criterion suggested by most psycholinguists, who tend to talk in informal, Gricean terms. While pragmatic theory can contribute to the development of an adequate criterion, it also stands to gain from the fact that disambiguation and reference assignment are more amenable to experimental testing than the recovery of implicatures. Here, collaboration between pragmatists and psycholinguists should be of benefit to both.

Robyn Carston has studied the contribution of enrichment processes to explicatures in a series of important papers;¹⁶ the role of inference in explicit communication is now being actively explored both inside and outside the relevance-theoretic framework.¹⁷ Criteria have been proposed for distinguishing explicatures from implicatures, and a case has been made for reanalysing some of Grice's best-known examples of generalised implicatures (e.g. the temporal implicatures carried by conjoined utterances, the quantity implicatures carried by numerals such as 'two' and 'three') as pragmatically determined aspects of explicit content. Much of this case rests on an intuitive distinction between truth-conditional and non-truth-conditional content which is standard throughout the speech-act and Gricean literature, but which could itself be usefully reassessed.

The claim that pragmatic principles can contribute to explicit content as well as implicatures has been seen as problematic by those who, following Gazdar (1979), thought of the semantics-pragmatics distinction in a rather non-Gricean way. Gazdar imported into pragmatics a picture common enough in formal semantics at the time, which conflated linguistic semantics with truth-conditional semantics and defined pragmatics as 'meaning minus truth conditions'. On this account, pragmatic processes should be 'post-semantic', and should not 'intrude' into the truth-conditional domain.

Relevance theorists have consistently rejected this picture.¹⁸ In *Relevance* (chapter 4, sections 1 and 7), following Fodor (1975), we systematically distinguished between linguistic semantics (the semantics

of natural-language sentences) and truth-conditional semantics (the semantics of conceptual representations). On this approach, the pragmatic processes that contribute to explicit truth-conditional content do not 'intrude' into a unitary semantics: they act on the output of linguistic semantics, enriching incomplete logical forms into fully propositional forms which are in turn the bearers of truth conditions. The need for some such distinction – which is not original to relevance theory – is now widely accepted by those working both inside and outside the relevance-theoretic framework.

2.2 Linguistic semantics

The implications of relevance theory for linguistic semantics, and in particular for what is traditionally regarded as non-truth-conditional linguistic meaning, have been a second major focus of research. In previous frameworks, non-truth-conditional meaning was typically analysed in speech-act terms. Speech-act semanticists treated a range of non-truth-conditional expressions (mood indicators, discourse adverbials, discourse particles, parentheticals) as indicators of illocutionary force. Grice extended this account to a range of non-truth-conditional discourse connectives, which he treated as conventionally implicating the performance of higher-order illocutionary acts.¹⁹ Within the relevance-theoretic framework, this approach to non-truth-conditional meaning is being reassessed.²⁰

Much of this reassessment was inspired by Diane Blakemore (1987), who reanalysed Grice's discourse connectives using a distinction between conceptual and procedural encoding; her account of discourse connectives as encoding procedural constraints on implicatures has provoked a flood of research.²¹ A further impetus was provided by our arguments against speech-act accounts of mood indicators in Wilson and Sperber (1988a), and by our more general critique of speech-act theory in *Relevance* (chapter 4, section 10).²²

In Wilson and Sperber (1993), we argued that mood indicators and discourse particles are best analysed in procedural rather than conceptual terms. In the relevance-theoretic framework, both types of expression contribute to explicatures rather than implicatures. We therefore generalised Blakemore's notion of constraints on implicatures, arguing that procedural meaning can constrain any aspect of the inferential phase of comprehension, whether explicit or implicit. We also questioned the assumption that procedural meaning and non-truth-conditional meaning invariably coincide: some expressions (e.g. discourse adverbials) which are standardly treated as non-truth-conditional may be best seen as encoding concepts; some truth-conditional expressions (e.g. pronouns)

may be best seen as encoding procedures. Relevance-theoretic alternatives to speech-act accounts of mood indicators, discourse particles, discourse adverbials and parentheticals sketched in that paper are now being actively explored.²³ It may turn out that the conceptual-procedural distinction will shed more light on linguistic semantics than the traditional distinction between truth-conditional and non-truth-conditional meaning.

2.3 Interpretive dimensions of language use

More fundamental than any of the distinctions discussed above is the one drawn in *Relevance* (chapter 4, sections 7–9), between descriptive and interpretive dimensions of language use. We claimed that, on the most basic level, every utterance is a more or less faithful interpretation of a thought the speaker wants to communicate. An utterance is descriptively used when the thought interpreted is itself entertained as a true description of a state of affairs; it is interpretively used when the thought interpreted is entertained as an interpretation of some further thought: say, an attributed or a relevant thought. In the light of this distinction, traditional pragmatic categories, e.g. tropes and speech acts, must be radically rethought: for example, metaphor falls together with descriptive uses of language, while irony, interrogatives and exclamation fall together as varieties of interpretive use.

Our approach to metaphor and irony, developed in a series of later papers, has been extensively discussed.²⁴ Perhaps surprisingly, most reactions have come not from Gricean pragmatists, whose analyses we severely criticised, but from psychologists, non-Gricean pragmatists and literary theorists. The range of data now being considered, and the range of explanations on offer, are much richer than those discussed in the rather limited Gricean literature.

The interpretive dimension of language use is not restricted to irony. Translation has been reanalysed from this perspective in a series of interesting works by Ernst-August Gutt.²⁵ The notion of interpretive use has also shed light on a range of traditional linguistic topics such as interrogatives, exclamation, echo questions, pseudo-imperatives, hearsay particles and metalinguistic negation, most of which have resisted analysis in purely descriptive terms.²⁶ There is much more to be done in this area, from both descriptive and theoretical points of view. However, the reorganisation proposed in *Relevance* seems to be bearing fruit.

2.4 Wider domains

A start has been made on investigating the implications of relevance theory in wider domains. In literary studies, the suggestions made by

Paul Kiparsky (1987) have been actively pursued.²⁷ Humour, politeness, advertising, argumentation, political language and language in education have all been investigated from a relevance-theoretic perspective.²⁸ Ruth Kempson has applied the assumptions of relevance theory to the investigation of generative grammar and issues of linguistic modularity.²⁹ Foster-Cohen (1994) and Watson (1995) have looked at language development; the broader implications of relevance theory for language acquisition are assessed in Smith (1989), Smith and Tsimpli (1995); theoretical considerations bearing on both evolution and development are discussed in Sperber (1994a).

In psychology, interesting results are being obtained in several domains. Frith (1989) and Happé (1991, 1992, 1993) have applied relevance theory to the analysis of autism. Politzer (1993) has reanalysed several major experimental paradigms in the psychology of reasoning, and shown how considerations of relevance affect the performance of subjects in ways that can explain some of the most striking experimental results. Sperber, Cara and Girotto (forthcoming) have reanalysed the literature on Wason's famous Selection Task, where subjects are asked to select evidence potentially relevant to evaluating the truth of a conditional statement. Sperber *et al.* suggest that the performance of subjects can be explained on the basis of intuitions of relevance developed in the process of comprehending the task. Their analysis yields precise and novel predictions involving the manipulation of effect and effort, which have been experimentally confirmed.

3 Revisions

3.1 *Not one but two Principles of Relevance*

In *Relevance*, we make two fundamental claims, one about cognition, the other about communication:

- (1) Human cognition tends to be geared to the maximisation of relevance.
- (2) Every act of ostensive communication communicates a presumption of its own optimal relevance.

Claim (2) is what we called the Principle of Relevance. However, many readers, even careful ones, have used the term 'Principle of Relevance' to refer to claim (1). This is a straight misreading, but an understandable one. Claim (1) is more fundamental and general than claim (2), and at least as worthy to be called a principle. We originally called claim (2) a principle to contrast it with other pragmatic 'principles' proposed in the literature: in particular Grice's Co-operative Principle. We failed to foresee that

when our book was read and interpreted – as we wanted – in the context of wider cognitive concerns, this use of the term 'principle' would seem rather arbitrary, cause unnecessary effort, and hence (as we should have predicted on relevance-theoretic grounds) lead to misinterpretation.

We have decided to remedy the situation by talking in future of two Principles of Relevance: the First (or Cognitive) Principle is given in (1), and the Second (or Communicative) Principle is given in (2). Throughout this book, the term 'Principle of Relevance' refers to the Second, Communicative Principle. The change is, of course, expository and not substantive, but it is worth spelling out what we hope to highlight by this reformulation.

3.2 *The First Principle of Relevance*

The First Principle of Relevance is less subtle than the Second Principle, but it is still controversial and in need of justification. As stated, it is also too vague, and in need of elaboration.

Relevance is not a commodity; it is a property. What is it a property of? By our definition, it is a property of inputs to cognitive processes. It can be a property of stimuli, for example, which are inputs to perceptual processes, or of assumptions, which are inputs to inferential processes. Stimuli, and more generally phenomena, are found in the environment external to the organism; assumptions, which are the output of cognitive processes of perception, recall, imagination or inference, are internal to the organism. When we claim that human cognition tends to be geared to the maximisation of relevance, we mean that cognitive resources tend to be allocated to the processing of the most relevant inputs available, whether from internal or external sources. In other words, human cognition tends to be geared to the maximisation of the cumulative relevance of the inputs it processes. It does this not by pursuing a long-term policy based on computation of the cumulative relevance achieved over time, but by local arbitrations, aimed at incremental gains, between simultaneously available inputs competing for immediately available resources.

Why assume that human cognition tends to be geared to the maximisation of relevance? The answer comes in two stages, one to do with the design of biological mechanisms in general, the other with efficiency in cognitive mechanisms.

We start from the assumption that cognition is a biological function, and that cognitive mechanisms are, in general, adaptations. As such, they are the result of a process of Darwinian natural selection (although other evolutionary forces may have helped to shape them). We assume, then, that cognitive mechanisms have evolved in small incremental steps,

mostly consisting in the selection of a variant that performed better at the time than other variants that were around. There are many ways in which one variant of a biological mechanism can perform better than others. There may be a qualitative difference in the type of benefits that different variants produce; or the difference may be quantitative, as when the same benefit can be achieved to a greater degree, or at a lower energy cost.

Whereas selection pressures for qualitative improvements vary perpetually with changes in the genotype and the environment, selection pressures for quantitative improvements are a relatively stable factor. *Ceteris paribus*, greater benefits or lower costs are always a good thing. In principle, there are many equally satisfactory ways of balancing costs and benefits: many ways, that is, of being efficient (although few, if any, may be genuine alternatives at a given point in the evolution of an adaptation). Hence, it is not possible to predict what exact balance of costs and benefits should be achieved in a given biological mechanism as a result of the pressure towards greater efficiency. What we can expect is that, in general, an enduring biological mechanism with a stable function will have evolved towards a better cost-benefit balance, i.e. towards greater efficiency.

For example, we can expect that the structure, placement and mode of operation of a muscle will *tend* to minimise the energy costs of performing the bodily movement it is its function to produce. Similarly, we can expect to find a *tendency* towards maximal efficiency in the design of cognitive mechanisms.

We assume, too, that human cognition is the joint product of many specialised mechanisms (see Barkow, Cosmides and Tooby 1992; Hirschfeld and Gelman 1994). Each cognitive mechanism contributes its qualitatively different benefits, in the form of cognitive effects. For each, there has been pressure towards cost-benefit optimisation.

All these cognitive mechanisms taken together constitute the cognitive system. The efficiency of the cognitive system as a whole depends on how its various sub-mechanisms are articulated with one another, and how the resources of the system are shared among them. Articulation and allocation of resources must be such as to maximise the likelihood that the most relevant available information will be processed in the most relevant way.

What the First Principle of Relevance says is that human cognition *tends* to be organised so as to maximise relevance. There may be many shortcomings, many cognitive sub-mechanisms that fail to deliver enough effect for the effort they require, many occasions when the system's resources are poorly allocated. The First Principle does not rule these out. Still, for it to be of any use, the tendency towards

maximisation of relevance must be strong enough overall to help guide human interaction. After all, The Second, Communicative Principle of Relevance is grounded in the First Principle, and in the further assumption that the First Principle does indeed make the cognitive behaviour of another human predictable enough to guide communication.

3.2.1 The First Principle of Relevance and truth Our definition of the relevance of an assumption in a context takes no account of the objective truth or falsity of the assumption itself, or of the conclusions that may be derived from it in the context. Thus, a false assumption that contextually implies many false conclusions, or a true assumption that combines with a false contextual premise to imply many false conclusions, is, by our definition, as relevant as a true assumption that implies many true conclusions. On the other hand, our rationale for introducing this notion of relevance has to do with considerations of cognitive efficiency, and the notion of cognitive efficiency cannot be divorced from that of truth. The function of a cognitive system is to deliver knowledge, not false beliefs. Does this mean there is something missing from our definition of relevance? Definitely, and it is in need of revision. Note, though, that for most of our purposes our incomplete definition is good enough.

When we use the notion of relevance to help describe how a cognitive system allocates its resources, there is no harm in leaving objective truth or falsity out of account. The system has no other way of distinguishing true from false assumptions than via its own inputs and internal processes. Basically, if an assumption is caused by the environment in the appropriate way (e.g. through perception), the system accepts it; if an assumption is inferentially derived by the system's own computational mechanisms from accepted premises, it again accepts it. When the system is a reflective one, e.g. a human being, it may be aware that it wants real knowledge and not false beliefs; it may be aware of the risk of accepting false assumptions; it may develop some procedures to double-check the outcome of other procedures; but all it can do in the end is trust the sum of its own procedures to deliver knowledge. So the system will take the output of its own mechanisms as cognitively warranted, and will assess relevance in terms of all contextual effects achieved, even though, unbeknownst to it, some of its conclusions may turn out to be false. From this solipsistic point of view (in the sense of Fodor 1980), truth can safely be ignored.

However, this is not the only point of view that needs to be taken into account. A reflective cognitive system may be aware that some of its beliefs are likely to be false, even if it cannot tell which, and it may regard information leading to false beliefs as worse than irrelevant.

Similarly, a reflective cognitive system that communicates with other systems may regard only true information as relevant to them. Take a speaker who wants her audience to think she is married, when in fact she is not. She lies:

(3) I am married.

Does she believe that what she says is relevant to the hearer, or only that it may *seem* relevant to him, since it would have been relevant if true? We suggest the latter.

Relevant information is information worth^ehaving. False information is generally not worth having; it detracts from cognitive efficiency. How should we incorporate this epistemic feature into our definitions? There are two possibilities: we might say that inputs to cognitive processes are relevant only if they meet some specific epistemic condition; or we might say that inputs are relevant only if the outputs of their cognitive processing meet some specific epistemic condition.

The most obvious, and apparently simplest, solution is to make truth of the input a necessary condition of relevance. There are three problems with this choice. First, we want to attribute relevance not just to assumptions but also to phenomena, and in particular to ostensive stimuli. These are inputs to cognitive processes, but they are not the sort of things that can be true or false. Utterances, of course, are said to be true or false, and they are a kind of ostensive stimulus; but when we say that an utterance is true, we really mean that its interpretation is true, and this is the output of a cognitive process of comprehension.

Second, truth of the conclusions seems more crucial to relevance than truth of the premises. Consider the following scenarios:

- (4) Peter is a jealous husband. He overhears Mary say on the phone to someone, 'See you tomorrow at the usual place.' Peter guesses rightly that she is speaking to a man, and infers, quite wrongly, that she has a lover and does not love him any more.
- (5) Peter is a jealous husband. He overhears Mary say on the phone to someone, 'See you tomorrow at the usual place.' Peter guesses wrongly that she is talking to a man, and infers, rightly as it happens, that she has a lover and does not love him any more. (Mary's lover is a woman.)

In (4), Peter's assumption that Mary was talking to a man was true, and led to rich contextual effects. However, these effects were false beliefs. Was Peter's assumption relevant? We would rather say that it seemed relevant, but in fact was not. In (5), by contrast, Peter's assumption that Mary was talking to a man was false, but it led to many true beliefs, so that here we would be willing to say that it was genuinely relevant

(though perhaps not as relevant as it seemed, since it also led to some false beliefs).

Take the more general case of fiction. When you hear a parable, or read *War and Peace*, you may gain insight, through some form of analogical thinking, into yourself, your life, and the world as they are. If only true inputs were relevant, we would have to say such fictions were irrelevant. If truth of the output is what matters, then fictions can be relevant after all.

So let us explore the second way of amending our definition of relevance: by treating an input as relevant only if the output of its cognitive processing meets some specific condition. The basic idea is that for an input to be relevant, its processing must lead to cognitive gains. Now recall our strategy in the book. We first defined relevance in a context, and then relevance to an individual. Our definition of relevance in a context can be left unchanged. A context, even coupled with an inference engine, is not yet a cognitive system; it does not have a cognitive function, and does not stand to gain from true representations or lose by false ones. Relevance in a context is a formal property, interesting as such (with possible applications in Artificial Intelligence, for instance), and is best left as it is.

Things change when we move from relevance in a context to relevance to an individual (or more generally, to any cognitive system). Contextual effects in an individual are *cognitive effects* (a phrase we have used in articles written after 1986). They are changes in the individual's beliefs. An individual does stand to gain or lose by the truth or falsity of his beliefs, and he does have cognitive goals. An individual, were he to reflect on it, would not be interested in contextual effects *per se*, but only in so far as they contribute to his cognitive goals. This is easily built into our definition of relevance to an individual. Let us first define a *cognitive effect* as a contextual effect occurring in a cognitive system (e.g. an individual), and a *positive cognitive effect* as a cognitive effect that contributes positively to the fulfilment of cognitive functions or goals. Then we replace definitions (42) and (43) of chapter 3 with (6) and (7):

- (6) *Relevance to an individual* (classificatory)
An assumption is relevant to an individual at a given time if and only if it has some positive cognitive effect in one or more of the contexts accessible to him at that time.
- (7) *Relevance to an individual* (comparative)
Extent condition 1: An assumption is relevant to an individual to the extent that the positive cognitive effects achieved when it is optimally processed are large.
Extent condition 2: An assumption is relevant to an individual to the

extent that the effort required to achieve these positive cognitive effects is small.

Definitions (58) and (59) of the relevance of a phenomenon to an individual should also be modified accordingly.

These changes in the definition of relevance might seem to raise two questions. First, isn't the notion of a positive cognitive effect far too vague? Well, we could have been more specific and defined a positive cognitive effect as an epistemic improvement, i.e. an increase in knowledge. All the effects we are actually considering in this book are of this relatively well-defined epistemic kind. However, we want to leave open the possibility of taking into account, in the full picture, other possible contributions to cognitive functioning, involving, for instance, the reorganisation of existing knowledge, or the elaboration of rational desires. And, yes, the resulting definition of a positive cognitive effect is vague, but that is a problem not for relevance theory, but for cognitive psychology in general.

The second question that this redefinition of relevance to an individual might seem to raise is this. Doesn't the First Principle of Relevance then become vacuous? If human cognition tends to be geared to the maximisation of relevance, and if relevance is itself defined in terms of positive cognitive effects, aren't we just saying that human cognition tends to be geared towards the production of positive cognitive effects; which, surely, is a truism, and a vague one at that?

In fact, the First Principle is far from a truism. It makes two empirical claims: neither is self-evident, and the second is original to relevance theory. The First Principle might be false: human cognition might achieve a balance of positive versus negative cognitive effects just good enough to avoid being selected out. In fact, human cognition, being an evolved and adapted system, reflects in fine-grained aspects of its design repeated past pressures towards optimisation. Moreover, we claim that there is one general and essential way in which human cognition exhibits good design, and that is by tending to allocate its resources to the processing of available inputs in such a way as to maximise the expected cognitive effects. That said, we ourselves have stressed that what we now call the First Principle of Relevance is indeed vague and general, and that what makes it worth stating are some of its precise and non-trivial consequences: in particular, the Second Principle of Relevance.

3.3 *Revising the presumption of relevance*

The (Second) Principle of Relevance states that every act of ostensive communication communicates a presumption of its own optimal

relevance. The presumption of relevance itself was spelled out as follows:

(8) *Presumption of optimal relevance*

- (a) The set of assumptions *I* which the communicator intends to make manifest to the addressee is relevant enough to make it worth the addressee's while to process the ostensive stimulus;
- (b) The ostensive stimulus is the most relevant one the communicator could have used to communicate *I*.

We believe that this formulation should be substantively modified. The modifications will make the presumption of relevance simpler, and we will argue that they not only preserve the predictive power of the earlier version, but significantly increase it.

There are two reasons for crediting a communicator with the intention to convey a presumption of relevance; these are reflected in the two clauses of the presumption. First, the communicator must intend her ostensive stimulus to appear relevant enough to the addressee to be worth his attention. Otherwise, he might not pay it enough attention, and communication would fail. This sets a lower limit on the level of relevance the communicator intends the addressee to expect. A version of this idea is built into clause (a) of the presumption of relevance above. In this version, the level of effort needed to reconstruct the intended interpretation is treated as given, and the presumption is that the effect will be high enough for the overall relevance of the stimulus to be at or above the lower limit (below which the stimulus would not be worth processing). Clause (a) says, in essence, that the level of effect is at least sufficient.

Now suppose we treat the level of effect rather than effort as given. Then by the same reasoning – based on the fact that the communicator must intend her ostensive stimulus to appear relevant enough – the addressee can have legitimate expectations about the level of effort needed to achieve this effect. This level of effort must be low enough for the overall relevance of the stimulus to be at or above the lower limit.

Since there is no principled asymmetry here between effect and effort, clause (a) of the presumption of relevance can be made both simpler and more general, as follows:

- (9) The ostensive stimulus is relevant enough for it to be worth the addressee's effort to process it.

Is it ever legitimate for the addressee to expect – and the communicator to intend him to expect – a level of relevance that is not merely at but well above the lower limit? Grice and most of his followers suggest that it is. They assume that speaker and hearer must have a common

goal that goes beyond merely understanding and being understood, and are expected to provide whatever information would best further this common goal. What is to be expected is not just relevance enough, but maximal relevance to achieving the common goal.³⁰

We have expressed disagreement with this view. It may be true that in most verbal exchanges the participants share a purpose that goes beyond merely understanding one another, but it need not always be the case. Conflictual or non-reciprocal communication, for example, involve no such purpose. It is also true that understanding is made easier by the presence of a common goal. We can account for by this pointing out that a common goal creates a number of mutually manifest contextual assumptions on which the interlocutors can draw. The existence of a common conversational goal need not be built into pragmatic principles. We still believe this is correct.

However, we ourselves have stressed that interlocutors always share at least one common goal, that of understanding and being understood. It is in the communicator's manifest interest both to do her best and to appear to be doing her best to achieve this common goal. This provides a second reason for crediting her with the intention to convey a presumption of relevance, and is reflected in clause (b) of the presumption as stated above. In its current version, however, clause (b) is wholly about effort. The intended effect is treated as given, and clause (b) says that the stimulus used to achieve this effect is the one that requires least effort from the addressee.

The presumption of minimal effort expressed by (b) is at best too vague and at worst too strong. A communicator may well be willing to try to minimise the addressee's effort, since this will make him more likely to attend to her ostensive stimulus and succeed in understanding it. Still, for all sorts of reasons, the particular stimulus she produces may not be the one that would absolutely minimise the addressee's effort. In the first place, there is the communicator's own effort to consider. As speakers, we are prepared to make only so much effort in formulating our thoughts, and as hearers, we know better than to expect flawlessly crafted utterances. Then there may be rules of etiquette or standards of ideological correctness that rule out the utterance that would be easiest to process (which would also be likely to convey unwanted weak implicatures). As speakers, we avoid what we see as objectionable formulations, and as hearers, we expect such restraint.

Clause (b) of the presumption of optimal relevance should in any case have allowed for the speaker's right to be lazy or prudish, i.e. to have her own preferences and take them into account.³¹ In later publications or oral presentations, we amended this effort clause to say that no *unjustified* or *gratuitous* effort was to be demanded. In other words,

from a range of possible stimuli which were equally capable of communicating the intended interpretation and equally acceptable to the communicator (given both her desire to minimise her own effort and her own moral, prudential, or aesthetic preferences), the communicator should prefer, and appear to prefer, the stimulus that would minimise the addressee's effort.

However, this line of reasoning, which was based on considerations of effort, applies equally to the effect side. Suppose that, from the communicator's point of view, her goals would be equally well served by a number of utterances (or other stimuli), all of which would cause the intended contextual effects, but some of which would cause further contextual effects, and be (or seem) more relevant to the addressee as a result. Which should she choose? She should choose the utterance that would be (or seem) most relevant to the addressee, for just the reasons given above in discussing the minimisation of effort.

Here is an illustration. Mary wants to make it quite manifest to Peter that she will be out from 4 o'clock to 6 o'clock. She might inform him of this by saying any of (10a-c):

- (10) (a) I'll be out from 4 to 6.
- (b) I'll be out at the Jones's from 4 to 6.
- (c) I'll be out at the Jones's from 4 to 6 to discuss the next meeting.

Suppose she assumes that any of these utterances would be relevant enough to Peter. Suppose it doesn't matter to her whether she tells him where she is going and why. Suppose the amount of effort needed to produce any of these utterances makes no difference to her. Then it would be rational enough to utter any of (10a-c), since each would achieve her goal at an equally acceptable cost to her. However, it would be most rational to produce the utterance most relevant to Peter, since this would make it most likely that he would attend to her communication, remember it, and so on: in other words, it would maximise the manifestness to Peter of the information that Mary wants him to have. Since (10c) would demand more effort from Peter than (10b), and (10b) than (10a), Mary should choose one of these longer utterances if and only if the extra information conveyed yields enough effect to make it more relevant to Peter. If he doesn't care where she is going, she should choose (10a). If he cares where she is going, but not why, she should choose (10b). If he cares both where and why, she should choose (10c). These choices are rational even if Mary doesn't particularly want to be helpful to Peter by telling him what he may want to know. They are rational as ways of maximising the chances that she will succeed in making manifest to him the one thing she does want to make manifest: that she will be away from 4 o'clock to 6 o'clock.

We can thus make the following generalisation. Take a set of stimuli that meet the following conditions: any of them would be likely to communicate what the communicator wants to communicate; she is capable of producing any of them; and she has no preferences among them, apart from wanting to choose the one that will be most effective in achieving her communicative goal. These stimuli may differ in terms of the effort demanded of the addressee, the effects achieved, or both effect and effort. The communicator should choose the stimulus that appears most relevant to the addressee, since this will make her communication most likely to succeed.⁵ For the same reason, she should appear to be choosing the stimulus that is most relevant to the addressee. In normal conditions, appearance and reality are likely to coincide.

The communicator's choice of ostensive stimuli is limited not only by her preferences but by her abilities. On the effort side, there may be stimuli that would be easier for the hearer to process, but that the communicator is unable to think of at the time, as when the best formulation of some thought just fails to come to mind. On the effect side, the limits on the communicator's abilities are even more significant. There may always be information that the hearer would find more relevant than anything the communicator has to offer. She cannot be more relevant than her own knowledge permits. If she decides to communicate in bad faith, and tries to make manifest assumptions that she does not believe, she would still want the addressee to think that what she is trying to communicate is warranted by what she knows.

Again, there is no principled asymmetry between effect and effort. The presumption is that, of all the stimuli that are available to her and acceptable as a means of achieving her particular communicative goal, the communicator will choose one that is as relevant as possible to the addressee. The second clause of the presumption of relevance can be made both simpler and more general, as follows:

- (11) The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences.

We now have a fully revised presumption of optimal relevance:

- (12) *Presumption of optimal relevance (revised)*
 (a) The ostensive stimulus is relevant enough for it to be worth the addressee's effort to process it.
 (b) The ostensive stimulus is the most relevant one compatible with the communicator's abilities and preferences.

This says that the addressee is entitled to expect a level of relevance high enough to warrant his attending to the stimulus, and which is, moreover,

the highest level of relevance that the communicator was capable of achieving given her means and goals.

3.3.1 *The Second Principle of Relevance: that the presumption of optimal relevance is ostensively communicated* It would be a mistake to read the presumption of optimal relevance, in either the early or the revised version, as describing a goal that rational communicators should achieve. Unlike Grice's maxims, neither the principle nor the presumption of relevance is presented as a goal to be pursued or a rule to be followed by the communicator. The (Second) Principle of Relevance is a descriptive (as opposed to normative) claim about the content of a given act of ostensive communication. It claims that part of that content is a presumption that this very act of communication is relevant to the addressee.

The addressee's aim in interpreting an utterance is to identify the communicator's informative intention. As with any attribution of an intention to an agent, this is done by observing the means she chooses and assuming that these are appropriate to her goals, given her beliefs. We claim that a presumption of optimal relevance is communicated by any act of ostensive communication. Given our definition of ostensive communication, for this to be true it must be mutually manifest to communicator and addressee that the communicator has the informative intention of making the presumption of relevance mutually manifest. We will now show that this is so.

A rational communicator must intend the stimulus she uses to appear relevant enough to the addressee to attract his attention and make him willing to spend the effort needed for comprehension. How relevant is that? There is a limit below which the addressee will be unlikely to attend to the stimulus at all; clearly, the communicator must intend the addressee to expect a level of relevance at least as high as this. Moreover, it is to the communicator's advantage that the addressee should expect a level of relevance well above this lower limit, so that he will be willing to invest the effort needed for comprehension. However, just as the addressee is guided in interpreting the utterance by the assumption that the communicator is rational, so the communicator's intentions are constrained by the assumption that the addressee is rational. A rational addressee will not expect more relevance than the communicator is willing and able to achieve. There is no point in expecting the communicator to give information she doesn't have, or to produce stimuli she is unable to think of at the time. Nor can she be expected to go against her own preferences. So a rational communicator intends her stimulus to appear as relevant as is compatible with her abilities and preferences.

In other words, it is necessary for the first clause of the presumption

of relevance to be manifest to the addressee, and it is advantageous for the second clause to be manifest too. A rational communicator should therefore want both clauses of the presumption of relevance to be manifest. We claim that this is not some hidden fact about the psychology of communicators, but is manifest to any competent communicator or addressee. Thus, when a communicator makes it mutually manifest to herself and her addressee that she is trying to communicate by means of a given stimulus, she thereby makes it mutually manifest that she intends a presumption of relevance to be manifest. Given our definition of ostensive communication, this amounts to saying that a presumption of relevance is communicated.

3.3.2 Some consequences of the revised presumption of relevance All the analyses we have given in this book and elsewhere on the basis of the old presumption of optimal relevance go through as before. It is still true that the rational way to go about interpreting an utterance, or any other ostensive stimulus, is to follow a path of least effort and stop at the first interpretation that satisfies one's expectation of relevance. However, in the old version, the expected level of relevance was systematically at the lower limit. This did not mean that an utterance could never be more than just relevant enough to be worth the hearer's attention. What it did mean is that in order to achieve a higher level of relevance, the speaker had to formulate her utterance so that the first interpretation that was relevant enough to be worth the hearer's attention would actually be more than relevant enough.

To illustrate, suppose that Mary says to Peter:

- (13) You remember I bought that lottery ticket? Well, guess what? I won £10,000!

Mary's statement, taken literally, may well be not only relevant enough to be worth Peter's attention, but much more relevant than he would have expected, given the unrevised presumption of relevance. Still, if this is the first accessible interpretation that is relevant enough (and unless it conflicts with other of his contextual assumptions), he will accept it as the one intended. This, at least, is what an analysis based on the unrevised presumption of relevance would (correctly) predict.

Compare this with the case where Mary says to Peter:

- (14) You remember I bought that lottery ticket? Well, guess what? I won a prize!

Here, the first accessible interpretation that is relevant enough will probably represent Mary's prize as just big enough to be worth talking about. If just knowing that she won a prize is relevant enough, then the

value of the prize may not be seen as relevant at all. Here again, an analysis based on the unrevised presumption of relevance is adequate.

The revised presumption of relevance yields the same analysis of these and similar examples. In interpreting (13), Peter assumes that Mary had the ability – in this case the knowledge – to say something more than minimally relevant (namely that she had won £10,000), and that she gave this information in the absence of contrary preferences. In interpreting (14), let us assume that Peter accepts clause (b) of the presumption of relevance and expects Mary's utterance to be the most relevant one compatible with her abilities and preferences. Still, he has no reason to think she has a more relevant piece of information that she is reluctant to share with him; so he will assume that the prize is merely big enough to be worth mentioning. Quite often, the lower limit mentioned in clause (a) of the (revised) presumption of relevance will coincide with the higher limit mentioned in clause (b). The speaker has something just relevant enough to be worth saying, and says it.

In some cases, though, the revised presumption yields different, and better, analyses. Here we will consider two. The first is adapted from Grice (1989: 32). Peter and Mary are planning a holiday in France. Peter has just said that it would be nice to visit their old acquaintance Gérard if it would not take them too far out of their way. The dialogue continues:

- (15) (a) *Peter*: Where does Gérard live?
(b) *Mary*: Somewhere in the South of France.

As Grice notes, Mary's answer implicates (16):

- (16) Mary does not know where in the South of France Gérard lives

This implicature is easily explained in terms of Grice's maxims. Mary's answer is less informative than the first maxim of Quantity ('Make your contribution as informative as is required') would suggest. 'This infringement [...] can be explained only by the supposition that [Mary] is aware that to be more informative would be to say something that infringed the second maxim of Quality, "Don't say what you lack evidence for"' (Grice 1989: 32–33).

In the unrevised version of relevance theory, we would have to explain this implicature by noting that, in the situation described, it would generally be mutually manifest that Mary is expected and willing to co-operate in planning the holiday in France. From this assumption, together with the fact that her reply is not relevant enough to answer Peter's question, it can be inferred that she does not know exactly where Gérard lives. Then not only is (16) manifest but, given

Mary's co-operativeness, it is mutually manifest that she should want (16) to be manifest. Hence (16) is a proper implicature.

This analysis acknowledges the presence, in this particular situation, of the kind of co-operativeness which Grice regards as there in principle in every conversation. We have argued that Gricean co-operativeness is neither always at work, nor always presumed to be at work. In circumstances where the speaker is not expected to be co-operative, implicatures of the type in (16) do not go through.

Suppose, for example, that it is mutually manifest that Mary is dead against visiting Gérard. Then her answer would not carry the implicature in (16). She may have no more precise information about Gérard's whereabouts, or she may have it but be reluctant to give it, and there is no telling which. Here, a strict Gricean would have to say that Mary is at least partially 'opting out' of the Co-operative Principle and the first maxim of Quantity. Just as we would have had to explain the Gricean implicature in (16) by adding the contextual assumption that the speaker is co-operative, so a Gricean would have to explain the absence of the implicature by adding the assumption that the speaker is unco-operative.

Notice, now, that the same dialogue could carry a different implicature. Suppose it is mutually manifest that Mary knows where Gérard lives. Then her answer in (15b) would implicate not (16) but (17):

(17) Mary is reluctant to say exactly where Gérard lives.

This raises a problem for the Gricean, since it violates both the Co-operative Principle and the first maxim of Quantity, and implicatures are supposed to arise only on the assumption that the Co-operative Principle is in force. With the unrevised version of the presumption of relevance, this example would have raised a problem for us too. Let us suppose that the information that Gérard lives in the South of France is relevant enough to be worth Peter's attention (even though it is less relevant than he would wish). Then, on our unrevised account, Peter should stop short of constructing the implicature in (17).

With the revised presumption of relevance, we can explain both standard Gricean implicatures such as (16), and non-Gricean implicatures such as (17), which are caused by, and express, a refusal to co-operate. In (15), if it is mutually manifest that Mary would like to be more specific about where Gérard lives, then her response, together with clause (b) of the revised presumption of relevance, will imply that she is unable to be more specific. If it is mutually manifest that this implication increases the relevance of her utterance, then it will be not just implied but implicated. On the other hand, if it is mutually manifest that Mary could have been more specific, then her response, together with clause (b) of the presumption of relevance, will imply that she is

unwilling to be more specific. Again, if it is mutually manifest that this implication increases the relevance of her utterance, it will be implicated.

Note that here we are making a subtle and non-obvious claim. We are claiming that if it is mutually manifest to communicator and audience that an assumption contextually implied by an utterance increases its overall relevance, then it is (in general) mutually manifest that the communicator intended this implication to be manifest. In other words, this assumption is communicated (as an implicature). This follows from clause (b) of the revised presumption of relevance, which states that the utterance is the most relevant one compatible with the communicator's abilities and preferences. If a mutually manifest implication of the utterance contributes to overall relevance, and thus helps to confirm the presumption of optimal relevance, the inference that the communicator intended it to play this role is sound. It is obvious that the communicator is able to implicate this assumption. There is evidence that she is willing to implicate it, since she has willingly chosen a form of utterance that manifestly carries this implication, which helps to confirm the presumption of relevance that she herself has communicated.

The claim that manifestly relevant implications can be treated as implicatures has one striking consequence. Sometimes, the addressee may justifiably attribute to the communicator an implicature that she never in fact intended to communicate. Sound though it may be, the inference from the mutually manifest fact that an implication is relevant to the conclusion that it is implicated (i.e. intentionally made manifest) is a non-demonstrative one, and it may on occasion be false. Consider a slightly different version of dialogue (15) above. It is mutually manifest to Mary and Peter that Mary is willing to give him all the relevant information she has:

(18) (a) *Peter*: You said you were in touch with Gérard. Where does he live?

(b) *Mary*: Somewhere in the South of France, I don't know exactly where.

In (18b), Mary says that she doesn't know exactly where Gérard lives. As it stands, this utterance, made without further explanation, contextually implies that she misinformed Peter when she claimed to be in touch with Gérard. She might not have intended to make this implication manifest, and *a fortiori* she might not have wanted to implicate it. However, unless she explicitly cancels the implicature (for instance, by explaining how it is that she doesn't know where Gérard lives despite being in touch with him), she will be taken to have implicitly admitted that she misinformed Peter. As this example shows, just as the choice of

words may commit a speaker to unwanted explicatures, so the contextual implications of an utterance may commit her to unwanted implicatures.

The second type of case where the revised presumption of relevance leads to better analyses has been much discussed in the literature under the label 'scalar implicatures'.³² Here is a typical example. In most situations, the utterance in (19) would implicate (20) or (21):

- (19) Some of our neighbours have pets.
- (20) Not all of our neighbours have pets.
- (21) The speaker doesn't know whether all her neighbours have pets.

These implicatures do not always go through, as witness (22) and (23):

- (22) Some of our neighbours certainly have pets; maybe they all do.
- (23) (a) *Peter*: Do some of your neighbours have cats, dogs, goldfish, that sort of thing?
- (b) *Mary*: Yes, some of our neighbour do have pets; in fact they all do.

At first blush, these facts are reasonably well explained in Gricean terms. A speaker who knew that *all* her neighbours have pets and who merely said, without the sort of qualifications in (22) and (23), that *some* of her neighbours have pets, would be giving less information than required by the first maxim of Quantity. To preserve the assumption that the speaker is obeying the Gricean maxims, the hearer must take her to implicate that she doesn't know whether all her neighbours have pets, or more strongly, that not all them do.

This Gricean account is not without weaknesses. It leaves open the question of how much information is required on a given occasion by the first maxim of Quantity – and hence of when 'some' actually carries an implicature. Nor does it offer any obvious way of deciding when 'some' implicates 'not all' (which it seems to do most of the time), and when it merely implicates ignorance on the part of the speaker. Still, 'some' conveys 'not all' so often that the implicature from one to the other is considered by most Griceans (e.g. Levinson 1987) to be a case of 'generalised implicature', working as a default inference automatically made, though defeasible in the presence of negative evidence.³³

In the unrevised version of relevance theory, examples like (19) raise the following problem. Consider a situation where the fact that (at least) some of the speaker's neighbours have pets would be relevant enough to be worth the hearer's attention. Then, having recovered this basic interpretation (on which 'some' is compatible with 'all'), the hearer would have no reason to go further and assume that the speaker meant 'some, but not all'. This is not an altogether undesirable result. In some cases it makes the right prediction, as in the following dialogue:

- (24) (a) *Henry*: If you or some of your neighbours have pets, you shouldn't use this pesticide in your garden.
- (b) *Mary*: Thanks. We don't have pets, but some of our neighbours certainly do.

Here, it seems to us, the fact that at least some of Mary's neighbours have pets is relevant enough, and there is no reason to assume she meant that not all of them do (or that she doesn't know whether all of them do). Griceans who treat the inference from 'some' to 'not all' as a generalised implicature would have to claim that Mary's utterance does have this implicature, or that the hearer of (24b) would first make this inference and then (for what reason?) cancel it. Neither hypothesis seems plausible to us.

However, in some cases the predictions of the unrevised presumption of relevance are not obviously correct. This happens when the basic interpretation of 'some' (where 'some' is compatible with 'all') is relevant enough to be worth the hearer's attention, but when it would clearly be more relevant to the hearer to know whether 'not all' is the case too. An example is (25):

- (25) (a) *Henry*: Do all, or at least some, of your neighbours have pets?
- (b) *Mary*: Some of them do.

Here, Henry has made manifest that it would be relevant to him to know not only whether some of Mary's neighbours have pets, but whether all of them do. An unrevised relevance model, applied mechanically to this case, would predict that Henry should stop at the first interpretation that is relevant enough; this is clearly the one on which Mary is taken to communicate that she has at least some neighbours who have pets, and nothing more. This prediction is manifestly wrong. Mary's answer would normally be taken to convey that not all of her neighbours have pets.

It would, of course, be easy enough to apply the relevance model flexibly: one might argue, for instance, that someone who asks a question automatically makes it manifest that what he would consider relevant enough is nothing less than a full answer to his question, or an utterance at least as relevant as that. In that case, Mary's answer in (25b), understood as conveying only that she has at least some neighbours who have pets, would not be relevant enough. Standard relevance considerations would cause it to be interpreted as implicating³⁴ that not all her neighbours have pets, thereby satisfying Peter's expectation of adequate relevance.

However, we much prefer a model that can be applied mechanically. Isn't this what taking cognitive science seriously is all about? The revised

relevance model is much more satisfactory in this respect (by which we mean not that we have a full-fledged, mechanically applicable model, but that at least we don't need to invoke special factors, however plausible, to account for not-so-special cases). With the revised presumption of relevance, the analysis of example (25) goes as follows. Mary's answer makes it manifest that she is either unable or unwilling to inform Peter that all her neighbours have pets. Either implication would increase the relevance of her utterance. In fact, in most circumstances Mary's answer will make it manifest that she is unable (rather than unwilling). This inability can in turn be explained in two ways: either she doesn't know whether all her neighbours have pets, or she knows that not all of them do. If one of these mutually incompatible assumptions is manifest enough, it will (in general) be mutually manifest that Mary intended it to be manifest, since it increases the relevance of her utterance and is compatible with her preferences. The resulting interpretation is the one consistent with the principle of relevance.

Mary's answer in (25) is a case where the speaker has deliberately chosen to express a less informative proposition when a closely related, equally accessible and more informative proposition would have demanded no more effort, either from Mary or from the hearer. All such cases have a similar analysis. If the more informative proposition would not have been more relevant, there is no implicature. If the more informative proposition would have been more relevant, the utterance will be taken to implicate either that the speaker is unwilling, or (more commonly) that she is unable to provide the more relevant information. In the latter case, the communicator's inability may be due either to her not knowing whether the more relevant information is true, or to her knowing it to be false. If either of these two possibilities is manifest and relevant, it will be treated as an implicature.

3.4 *Far too early to conclude*

There are many other aspects of relevance theory that we would like to see developed, and that we or others have begun working on in articles and unpublished lectures. Many involve local revisions of the version of the theory presented in this book. Some open new perspectives that may turn out to be more important in the general balance of the theory than the present revisions.

Experimental studies testing relevance-theoretic hypotheses have just begun, and we hope that they will lead to revisions, new insights, and, perhaps more important, new problems to investigate. Interesting applications of the theory to literary studies suggest that it might be of some relevance, more generally, in the study of various cultural productions.

Novel insights and new problems should come from the formal modelling of the theory, possibly with the use of spreading activation models which seem particularly well suited to representing, on the one hand, the role of accessibility, and, on the other, the way the system's computations can be guided on line by monitoring its efforts and effects. Two important and related domains have hardly been explored at all from a relevance-theoretic perspective: the theory has been developed from the point of view of the audience of communicative acts, and without taking into account the complex sociological factors richly studied by sociolinguistics. The cognitive processes at work in the communicator, and the social character and context of communication are, of course, essential to the wider picture, to the study of which we hope relevance theory can contribute, and from which it stands greatly to benefit.

We ourselves have been working on a revised and more detailed description of inferential comprehension, integrating in particular the processes involved in enrichment and the comprehension of loose talk or metaphor. This work will be presented in our forthcoming *Relevance and Meaning*.