

<http://spilplus.journals.ac.za/>

RELEVANCE THEORY

EXPLAINING VERBAL COMMUNICATION

**Melinda Sinclair
and
Walter K. Winckler
Department of General Linguistics
University of Stellenbosch**

SPIL PLUS 18

1991

ACKNOWLEDGEMENTS

We are happy to acknowledge our indebtedness to various people for helping to make the creation of this text possible:

- * to Rudie Botha - for running a department that abounds in intellectual stimulus, and for nudging the two of us along into a writing project that has proved to be effortful, but worth our while;
- * to the other members of the 1990 departmental reading group that set out to see what relevance theory was all about;
- * to Christine Anthonissen, Christo van der Merwe, Marianna Visser and Hester Waher (UCT), the other members of the splinter group that opted for a more sustained attempt to get to grips with relevance theory - for sharing with us their interest and enthusiasm, for freely exchanging ideas about the form and content of a possible text, and for commenting so helpfully on the earliest drafts of chapter 2;
- * to Cecile le Roux, our editor - for her invaluable scrutiny of the pre-final draft and for welcome hints towards its improvement;
- * to Hildegard van Zweel - for her frank and perceptive comments on the text, and for her moral support;
- * to Christine Smit - for cheerfully and efficiently toiling away to help turn an often messy manuscript into respectable text;
- * to Suzette Winckler - for ready, able, and imaginative assistance in formatting the text for the laser printer, and,
- * last but indubitably not least, to Marius Sinclair and Suzette Winckler, the spice of our respective lives - for their generosity and understanding in allowing us all the extra time this project took.

We know better, alas, than to hope that all our readers will find all of this text optimally relevant; and so, for the errors of omission and commission that beset this text despite our combined best efforts, we hasten to claim jointly the full and exclusive responsibility.

C O N T E N T S

1 INTRODUCTION	1
2 ASPECTS OF UTTERANCE INTERPRETATION:	
A SKETCH OF THE HEARER'S TASK	2
2.1 Introduction	2
2.2 Recovering the semantic representation(s) of the sentence(s) uttered	2
2.3 Recovering the explicit content	4
2.3.1 Introduction	4
2.3.2 Disambiguation	5
2.3.3 Reference assignment	5
2.3.4 Enrichment of vague terms	6
2.4 Recovering the implicit content	7
2.4.1 Introduction	7
2.4.2 Recovering the implicatures of ordinary assertions	7
2.4.3 Recovering stylistic effects	8
2.4.4 Recovering a metaphorical interpretation	8
2.4.5 Recovering an ironical interpretation	9
2.4.6 Recovering the illocutionary force	9
2.5 Summary	9
3 RELEVANCE	11
3.1 Introduction	11
3.2 Utterance interpretation and human cognition	11
3.3 Relevance and information processing	13
3.3.1 The general goal of human cognition: efficient information processing	13
3.3.2 Relevant information	14
3.3.3 Processing effort	18
3.3.4 A comparative definition of relevance	19
3.4 The principle of relevance	20
3.5 The presumption of optimal relevance	23
3.6 Consistency with the principle of relevance	24

4	HOW CONSISTENCY WITH THE PRINCIPLE OF RELEVANCE	
	EXPLAINS UTTERANCE INTERPRETATION	27
4.1	Introduction	27
4.2	The role of decoding in utterance interpretation	27
4.3	How the consistency criterion explains the recovery of the explicit content of utterances	33
4.3.1	Introduction	33
4.3.2	Disambiguation	37
4.3.3	Reference assignment	44
4.3.4	Enrichment of vague terms	47
4.3.5	Summary	48
4.4	How the consistency criterion explains the recovery of the implicatures of ordinary assertions	50
4.5	How the consistency criterion explains the recovery of stylistic effects	62
4.6	How the consistency criterion explains the recovery of a metaphorical interpretation	67
4.7	How the consistency criterion explains the recovery of an ironical interpretation	72
4.8	How the consistency criterion explains the recovery of the illocutionary force of utterances	76
4.9	Explaining utterance interpretation: the common pattern	84
5	IN CLOSING	88
	REFERENCES	95
	CONTENTS OF PREVIOUS ISSUES OF SPIL PLUS	98

Chapter 1

INTRODUCTION

Proposals by Dan Sperber and Deirdre Wilson concerning a theory of relevance have attracted increasing scholarly attention in the last ten years or so. This is evidenced by the number of publications in which these proposals are critically discussed and/or applied.¹ Today, Sperber and Wilson's theory of relevance can hardly be ignored by scholars interested in describing and explaining verbal communication.

The main source on Sperber and Wilson's theory of relevance is of course their book *Relevance. Communication and cognition* (1986). In addition to their book, there are also available (inter alia) the authors' own *Précis* (1987) and at least two "outline" articles in which they present their ideas.² Nevertheless, when we set about getting a grasp of the theory, we found it hard work in spite of the availability of such a variety of sources.³ We soon became convinced that there was a need for a guide to relevance theory adapted to the situation of readers like ourselves: people keen to gain an understanding of how the theory operates in specific instances, but unable to set aside more than strictly limited time for this purpose. The present text is the outcome of an attempt to write just such a guide.

Specifically, this text is intended to help its readers understand the relevance-theoretic analysis of utterance interpretation. The readers we have in mind are not interested, in the first place, in foundational issues: including the arguments by which Sperber and Wilson seek to justify their basic concepts and principles, or by which they defend them against conceivable or existing alternatives. We take it that our readers, without denying the importance of these issues, are interested primarily in how the various aspects of utterance interpretation are to be described and explained in terms of relevance theory. Our text accordingly focuses on the analysis of specific examples of verbal communication, and presents theoretical constructs as and when they become indispensable. The copious references to other literature, and especially (Sperber and Wilson 1986), which we give throughout will nevertheless enable any reader who might need to do so to pursue foundational and technical issues.

In addition to this introductory chapter, our text contains four chapters. In chapter 2, we offer a preliminary overview of the various aspects of utterance interpretation. In chapter 3, the principle of relevance - the heart of Sperber and Wilson's theory - is introduced and illustrated. In chapter 4 - in its turn the heart of our text - we show the principle of relevance in operation by presenting details of analyses of the various aspects of verbal communication.⁴ In chapter 5, we conclude by offering some comments on the question, "What is the merit of relevance theory as a new approach in the field of pragmatics?"

¹ Bibliographical particulars of such publications are given in chapter 5 below.

² These outline articles are listed in our bibliography as (Wilson and Sperber 1986; 1987a). Note that in the rest of the notes we will use the abbreviations "W&S" and "S&W" to refer to Wilson and Sperber, and Sperber and Wilson, respectively.

³ We were not alone in this; the experience was shared by the other linguists in a local reading group which in 1990 attempted to assimilate the primary literature mentioned above.

⁴ This division into chapters 2, 3, and 4 owes much to the textual organisation adopted in (W&S 1986; 1987a).

Chapter 2

ASPECTS OF UTTERANCE INTERPRETATION: A SKETCH OF THE HEARER'S TASK

2.1 Introduction

Consider the task facing a hearer who is aware that he has been addressed by a speaker, and that the speaker intends by her utterance to communicate some information to him.^{1,2} For the communication to succeed, the information which the speaker intends to convey has to be identified by the hearer. What must the hearer do in order to identify this information?³ Or, in other words, what must the hearer do in order to interpret the utterance correctly? Correctly interpreting an utterance is a complex process. The aim of this chapter is to review some of the aspects of the process of utterance interpretation. Below, this process will be taken to involve three major subtasks for the hearer: (i) recovering the semantic representation(s) of the sentence(s) uttered; (ii) recovering the explicit content conveyed by means of the utterance; (iii) recovering the implicit content conveyed by means of the utterance.⁴

2.2 Recovering the semantic representation(s) of the sentence(s) uttered

In normal verbal communication, the hearer of an utterance has to recover the semantic representation(s) of the sentence(s) uttered by the speaker. This is so, because in normal verbal communication the speaker chooses a particular utterance in part for its semantic content.⁵ If the utterance is expressed in the hearer's native language, he automatically analyses the utterance as a token of a particular linguistic structure, and automatically assigns to the utterance the semantic representation which the grammar of the language pairs

¹ We follow S&W's (1986:256, n.13) helpful convention of referring to the speaker or communicator as female and to the hearer, audience or addressee as male, "unless the context indicates otherwise". All such references are of course to be understood in a generic sense.

² S&W (1986:2, 57-60, 73-75) use *information* in a broad sense, including for example propositional content whether true or false, propositional attitudes such as doubt or regret, and also "impressions" (an impression is inherently vague and is what results when the communicator merely "stirs the thoughts of the audience in a certain direction" (p.60)). Their notion of 'information' will be concretised by the various examples of utterance interpretation discussed in the course of our text.

³ From now on, following S&W, we will freely use *recover* as an umbrella term meaning "identify (whether by decoding, by demonstrative inference or by non-demonstrative inference)".

⁴ The dividing line between explicit content and implicit content presented here is taken over from S&W (1986:181-183). Their notion of such a dividing line is explicated in 4.3 below. The use of their distinction between explicit and implicit content at this point of the exposition will make it easier later on to link up the various aspects of utterance interpretation with the processes responsible for them.

⁵ This point is discussed by S&W (1986:178). Example (2) below is due to them.

with that structure and so, less directly, with its phonetic form as well.^{6,7} The grammar of the language concerned determines the linguistic description of the sentence(s) uttered, and this linguistic description yields a range of semantic representations, one for every sense of the sentence uttered. (This latter point will be taken up in 2.3.1 below.) Suppose, for instance, that the speaker wants to communicate to the hearer the assumption that an individual named Peter picked three apples.⁸ One way for the speaker to do so is by uttering the following:

- (1) Peter picked three apples.

Why is this? Obviously, because the grammar of the language assigns to the underlying sentence a suitable semantic representation. The semantic representation is suitable in that it contains information from which the intended interpretation of the utterance can be – as it happens, fairly straightforwardly – recovered.⁹ Obviously, too, for the hearer wishing to recover the intended interpretation, this means that he must recover the semantic representation associated in this way with the sentence uttered.

Note, however, that utterances can also be used to convey information which is not in any way connected to the semantic content of the sentence uttered. For instance, suppose that someone is asked whether she still stammers, and that she replies with the following utterance:

- (2) Peter Piper picked a peck of pickled pepper.

The speaker's intention is clear: to inform the hearer that she no longer stammers. Yet the information which the speaker intends to convey is in no way related to the semantic content of her utterance. Rather, she intends this information to be somehow inferred from the acoustic properties of her utterance. Clearly, then, the semantic content of the utterance plays no part in determining the information which is conveyed in the situation just outlined. This can be further illustrated by considering that the same information can, in such a situation, be as effectively conveyed by another utterance:

⁶ Cf. S&W 1986:174-177 for a more detailed presentation of these and other points relating to the part played in the interpretation of an utterance by the recovery of the semantic representation(s) of the sentence(s) uttered. It is of course true that two or more, successive, sentences can underlie an utterance; for convenience, however, our formulation from here onwards will provide for a single underlying sentence only (unless the utterance under consideration happens clearly to involve two or more underlying sentences in succession).

⁷ In terms of the distinction being appealed to here, the particular linguistic structure is the "type" of which the phonetic stimulus is a "token". Tokens are "unique physical entities, located at a particular place in space and time" (Lyons 1977:14); as such, "tokens instantiate their type" (Lyons 1977:13); and tokens "are identified as tokens of the same type by virtue of their conformity to the type that they instantiate" (Lyons 1977:14).

⁸ Informally, S&W (1986:2) take *thoughts* to mean "conceptual representations (as opposed to sensory representations or emotional states)" and *assumptions* to mean "thoughts treated by the individual as representations of the actual world (as opposed to fictions, desires, or representations of representations)".

⁹ In example (1), the intended interpretation of the utterance is linked fairly directly to the information contained in the semantic representation of the underlying sentence. This is atypical, as will be illustrated in 2.3 and 2.4 below; typically, the intended interpretation of an utterance is very much underdetermined by the semantic representation of the underlying sentence.

(3) The sixth sheik's sixth sheep's sick.

Of course, the semantic content of (2) is utterly unlike that of (3), and yet the fact is that either of these two utterances can be used to convey the information that the speaker no longer stammers. This fact strikingly illustrates the general point being made here: some utterances can be used informatively even though their semantic content plays no part in their interpretation.

Admittedly, the use of (2) in the way outlined above is atypical as a use of language for the purpose of communicating information. Normally, one of the reasons why the speaker chooses to utter one sentence rather than another is the fact that this sentence has the semantic content which it has. The information which the speaker intends to convey by means of her utterance is somehow connected to the utterance's semantic content. For the hearer faced with the task of interpreting an utterance, this has the obvious implication that he will have to recover its semantic content. To an important extent, the semantic representation recovered by the hearer forms the starting point for the process of utterance interpretation; as such, therefore, this semantic representation contributes substantially to the outcome of the process.¹⁰

2.3 Recovering the explicit content

2.3.1 Introduction

The recovery of the semantic representation(s) of the sentence uttered falls short, in several ways, of a recovery of the speaker's intended meaning. For one thing, the semantic representation(s) recovered do not represent the explicit content of the utterance.¹¹ This discrepancy is due to at least two reasons. Firstly, in the case of semantically ambiguous utterances, the grammar of the language associates more than one semantic representation with the underlying sentence. Before he can go on to work out a full representation of the speaker's intended meaning, the hearer therefore has to make the appropriate choice among the set of recovered semantic representations. Secondly, the appropriate semantic representation of the sentence uttered is in any event only a partial representation of the proposition expressed by the utterance. For a complete representation of this proposition to be constructed, firstly, the referring expressions in the semantic representation need to have referents assigned to them and, secondly, any vague, semantically incomplete, expressions in the semantic representation need to be made more specific.¹²

¹⁰ In 4.2 below the role which, in terms of relevance theory, the recovered semantic representation plays in utterance interpretation will be considered in some detail.

¹¹ The exact nature of the "gap" between the intended semantic representation of a sentence, and the proposition to be recovered by the hearer from an utterance of the sentence, is discussed and illustrated in detail in 4.2 below.

¹² For a discussion and illustration of what is involved in recovering the proposition expressed by means of an utterance cf. S&W 1986:179-180; 183-193. We present and illustrate their views in some detail in 4.3 below.

Assigning a unique, complete, proposition to an utterance, then, involves three subtasks for the hearer: disambiguation, reference assignment, and enrichment. Each of these subtasks is briefly illustrated below.

2.3.2 Disambiguation

Consider the following utterance:

- (4) Peter's bat is too big.

The sentence underlying this utterance is ambiguous between the noun *bat* meaning roughly 'small mouselike flying mammal active at night' and the noun *bat* meaning roughly 'wooden instrument for hitting the ball in certain games'. For convenience, let us shorten the wording of these two senses to "flying mammal" and "hitting instrument", respectively. Now, because of the ambiguity between the 'flying mammal' sense and the 'hitting instrument' sense the hearer's decoding of (4) will yield two semantic representations, one for each of the two possible sentence meanings.

Utterances which are ambiguous in the way illustrated by (4) make it clear that the hearer's task involves more than just recovering the semantic representation(s) of the sentence uttered. If the hearer is interested in determining what information the speaker intends to convey by her utterance (4), one of the recovered semantic representations must be selected as the basis for the further processes of utterance interpretation. And if he is to succeed in recovering the information which the speaker intends to convey, the hearer will have to select the right semantic representation - that is, the one which the speaker has in mind.

2.3.3 Reference assignment

Successfully disambiguating an utterance still leaves underdetermined various aspects of the intended interpretation of the utterance. One of these aspects is the interpretation of referring expressions such as the proper name *Peter* or the pronoun *she*.

Consider again utterance (4), *Peter's bat is too big*. Unless the hearer can tell who the speaker means by *Peter*, he is not going to be able to recover the information which the speaker intends to convey to him. But to work out or recognise who the speaker means by *Peter*, the hearer needs to know more than just the grammar of English. As Sperber and Wilson (1986:10) put it, taking *Bill* and *Betsy* as examples of proper names:

"A generative grammar cannot determine who ... 'Bill' and 'Betsy' refer to ... It can only provide some very general indications. It might state, for example, that ... 'Bill' and 'Betsy' refer to people or other entities with those names ..."

In short, the hearer's problem in interpreting proper names is that, considered at the level of sentence meaning, they are "referentially indeterminate" (Sperber and Wilson 1986:10, 188, 193). It is up to the hearer to "assign" the intended "reference" (or "referent") at the level of utterance interpretation, and to be able to do so he needs information beyond the purely grammatical information embodied in semantic representations of sentences.

Pronouns too are referentially indeterminate, when considered at the level of sentence meaning. Here is an example:

(5) She has had breakfast.

A hearer wishing to recover the information which a speaker intends to convey by means of an utterance of sentence (5) will have to identify the intended referent of the pronoun *she*. But, again, a generative grammar cannot determine who a pronoun such as *she* "refers to. It can only provide some general indications" (Sperber and Wilson 1986:10). It might state for example, in connection with references to people, that *I* always refers to the speaker, *you* to the hearer and *she* to a female person who is neither the speaker nor the hearer. Again, to resolve the referential indeterminacies associated with pronouns, the hearer will need information beyond the purely grammatical.

2.3.4 Enrichment of vague terms

The semantic representation recovered by the hearer may contain terms which are vague, or imprecise, and therefore semantically incomplete.¹³ Consider for instance the genitive expression *Peter's bat* in (4) above. It may be used to refer to a bat owned by Peter, to a bat chosen by Peter, to a bat killed by Peter, and so on.¹⁴ Likewise incomplete is the meaning of the adjectival phrase *too big*. Of course the speaker wishes to convey the information that the bat is *too big for something*. If the hearer does not know what this something is, he does not fully understand what *too big* is being used to express. The process of assigning more definite meanings to vague, or imprecise, terms is called *enrichment* by Sperber and Wilson (1986:185). To illustrate the notion of enrichment, let us consider one more example:

(6) It will take some time to repair your watch.

Clearly, in using (6), a speaker does not intend to express the mere truism that watch-repairing is a process with a temporal duration. Rather, she intends to inform the hearer that the job in question is going to take longer than would usually be expected. So the hearer has to enrich the meaning of *some time* by drawing on his expectations of how long the watchmaker concerned usually needs for repairs - say, one week. The more precise those expectations are, the more precisely the hearer can enrich the meaning of *some time* - say, to 'definitely more than one week'.¹⁵

¹³ On the notion of 'semantic incompleteness' cf. e.g. S&W 1986:10, 188, 193.

¹⁴ For some discussion of the vagueness of genitive expressions cf. S&W 1986:188. On the variety of interpretations allowed by the genitive construction cf. also Quirk et al. 1972: paragraphs 4.94, 13.27 and 13.28.

¹⁵ For a more detailed presentation of this example cf. S&W 1986:189.

2.4 Recovering the implicit content

2.4.1 Introduction

Suppose the hearer of an utterance has successfully completed the various subtasks outlined above. By so doing, he has assigned a unique proposition to the utterance. That is, he has found out what information is being conveyed explicitly by the utterance. For instance in the case of (4), *Peter's bat is too big*, he has identified the person who is the intended referent of *Peter*, he has identified the intended sense of *bat* as 'instrument for hitting a ball', and he has identified the intended enriched meaning of *too big* as 'too big for successful batting'.

As is well known, however, part of the information conveyed by an utterance is conveyed implicitly. Accordingly a hearer, in order to understand an utterance fully, must also find out what its implicit content is. The notion of the 'implicit content' of utterances - also called their *implicatures* - covers, in fact, a wide range of aspects of meaning.

2.4.2 Recovering the implicatures of ordinary assertions

Let us start off with one of the best-known cases of utterances conveying implicit information, namely the implicatures of ordinary assertions.¹⁶ Let us assume a situation in which the assertion in (7b) is uttered as a response to the question in (7a).

- (7) (a) *John*: Do you want some coffee?
(b) *Mary*: Coffee would keep me awake.

The proposition expressed by the assertion in (7b) is about the effect which coffee would have on the speaker, Mary. Note, in particular, that it contains no information about Mary's desire or non-desire for coffee. Nevertheless, intuitively it is clear that a speaker who uses this assertion in this situation intends to convey either the information that she wants coffee, or the information that she does not want coffee. This information is part of the implicit content, or implicatures, of the utterance. So one of the things to be decided by the hearer, John, interpreting the assertion in (7b) is whether the speaker intends to convey the information that she wants coffee, or the information that she does not want coffee. In making his decision, the hearer will have to take situational factors into account.

Note incidentally that, like the interpretation of the assertion in (7b), the interpretation of the question in (7a) also involves the recovery of an implicature. In the kind of situation assumed above, the question in (7a) is normally intended and appropriately interpreted, not merely as a request for information about Mary's desire for coffee, but also as an offer by John to provide the coffee, should Mary want some. The information that John is offering to provide coffee is conveyed implicitly.

¹⁶ For a more detailed presentation of a relevance-theoretic account of implicatures cf. S&W 1986:193-200. We will take up that account in 4.4 below.

2.4.3 Recovering stylistic effects

Compare the respective interpretations of the following utterances:

- (8) (a) My childhood days are gone.
(b) My childhood days are gone, gone.

Both utterances are normally taken to convey at least the information contained in the proposition expressed by (8a). The repetition of the adjective *gone*, however, has the effect of making (8b) convey something more than this. Again, this additional information is being conveyed implicitly. The repetition of the adjective may convey, for example, that the speaker is moved by a torrent of memories of her childhood (memories which the hearer is invited to imagine for himself). The repetition of the adjective can be said to achieve a particular type of stylistic effect, namely what Sperber and Wilson refer to as a *poetic effect*. When interpreting an utterance such as (8b), therefore, the hearer must also recover the speaker's intended stylistic effects.¹⁷

2.4.4 Recovering a metaphorical interpretation

Sometimes speakers use an utterance without intending to convey its explicit content, as contained in the proposition expressed by the utterance. Consider for instance an utterance such as (9).¹⁸

- (9) This room is a pigsty.

Normally, such utterances are intended to be interpreted not literally, but *metaphorically*. For instance, normally a speaker who utters (9) will not intend to convey the information that the room referred to is a place where (biological) pigs live. Rather, the speaker will intend to communicate that the room shares certain offensive properties with a pigsty. To be more specific, (9) will be intended to convey the information that the room is very dirty and untidy. In short, then: one further part of the hearer's task of recovering the information which the speaker intends to convey by an utterance is to determine whether the utterance is intended to be interpreted metaphorically and, if so, what the intended metaphorical interpretation is.

Note that there is an important difference between the interpretation of metaphorical utterances like (9), and the interpretation of utterances like (2), *Peter Piper picked a peck of pickled pepper*, as outlined in 2.2. In both cases the proposition expressed by the utterance does not form part of the information conveyed by the utterance. In the case of (9), however, this proposition does play a role in determining the information conveyed; that is, the information contained in this proposition serves as the hearer's basis on which to "guess" what information the speaker intends to convey. As was explained in 2.2., the semantic content of (2) plays no role whatever in determining the information conveyed by (2).

¹⁷ For a relevance-theoretic analysis of stylistic effects, including poetic effects, cf. S&W 1986:217-224. We present this analysis in 4.5 below.

¹⁸ For a discussion of metaphor within a relevance-theoretic framework cf. S&W 1986: 231-237. We explicate their account in 4.6 below.

2.4.5 Recovering an ironical interpretation

Consider the following utterance:

(10) It's a lovely day for a picnic, indeed.

Consider too the intended meaning of (10) in the following situation: John has organised a picnic, telling Mary that it is a lovely day for a picnic. They go for the picnic, and it rains. At some point Mary utters (10). How is the hearer, John, to interpret (10)? In the situation sketched, clearly the speaker, Mary, intends to convey an attitude of scorn or ridicule towards the opinion that it's a lovely day for a picnic. That is, in this situation (10) must be interpreted ironically rather than literally. In sum: determining whether an utterance has an ironical interpretation, and, if so, determining also the content of this ironical interpretation, constitute yet another aspect of utterance interpretation.¹⁹

2.4.6 Recovering the illocutionary force

Another component of the hearer's task is illustrated by utterances such as the following:

(11) Clean up your room.

The salient point is that this utterance can be used with more than one illocutionary force.²⁰ At least the following two possible illocutionary forces will have to be considered by the hearer: (i) that (11) is intended as a request (or even an order) to the hearer to clean up his room, and (ii) that (11) is intended as advising the hearer to clean up his room. The utterance (11) in itself, however, does not explicitly convey which of these is the intended illocutionary force. In sum: determining what illocutionary force the speaker of an utterance has in mind forms part of the hearer's task of recovering the implicit content of the utterance.²¹

2.5 Summary

Utterance interpretation is a complex process, involving a variety of (sub)subtasks. The first major subtask is to recover the semantic representation(s) of the sentence(s) uttered. The second major subtask is to recover the explicit content conveyed by means of the utterance. It involves three subsubtasks: (i) disambiguating the utterance, (ii) assigning referents to the referring expressions that occur in the utterance, and (iii) enriching vague terms that occur in the utterance. The third major subtask is to recover the implicit content conveyed by means of the utterance. This too may involve a number of subsubtasks: (i) recovering the

¹⁹ For a discussion of irony within a relevance-theoretic framework cf. S&W 1986:237-243. Both of the examples considered in the text above have been taken over from that discussion. We present S&W's relevance-theoretic account of irony in 4.7 below.

²⁰ The illocutionary force of an utterance reflects the illocutionary act(s) - for example, stating, offering, promising, requesting - which the speaker of the utterance can intend to perform by means of it.

²¹ For a discussion of speech acts within a relevance-theoretic framework cf. S&W 1986:243-254. We present their account in 4.8 below.

implicatures in the case of an ordinary assertion, (ii) recovering any possible stylistic effects, including poetic effects, (iii) recovering a possible metaphorical interpretation, (iv) recovering a possible ironical interpretation, and (v) recovering the illocutionary force of the utterance.²²

²² Conceptually, the distinction of 'task', 'subtask' and 'subsubtask' is important. Terminologically, however, we will not in the rest of the text consistently distinguish among the terms *task*, *subtask*, and *subsubtask*. The intended reference in each instance will be clear enough from the context.

Chapter 3

RELEVANCE

3.1 Introduction

The main aim of chapter 3 is to introduce the reader to the principle of relevance, the central principle of relevance theory. In 3.2 we present Sperber and Wilson's idea that general purpose processes of human cognition play a major role in utterance interpretation, and we explain the relationship between these general purpose processes and relevance theory. In 3.3 Sperber and Wilson's notion of 'relevance' is presented and illustrated. In 3.4 the principle of relevance itself is presented. In 3.5 we present a conception which is fundamental to this principle, namely 'the presumption of optimal relevance'. And, finally, in 3.6 we present the crucial notion of 'consistency with the principle of relevance'. As will be shown in chapter 4, it is this last notion that forms the basis for Sperber and Wilson's relevance-theoretic analysis of utterance interpretation.

3.2 Utterance interpretation and human cognition

To interpret an utterance, it has been suggested in chapter 2, is for the hearer to recover the information which the speaker intends to convey by means of the utterance. Moreover the hearer's task of recovering this information, it was also suggested, involves a number of conceptually distinct subtasks. Equivalently, the recovery of this information may be looked upon as involving a corresponding number of conceptually distinct processes.

How do the processes at work in utterance interpretation relate to human cognition in general? Sperber and Wilson suggest an answer to this question within the framework of their "modular" view of the structure of the human mind. They have developed this view of theirs on the basis of an approach presented by J.A. Fodor in his 1983 book *The modularity of mind* (Cambridge, Massachusetts: The M.I.T. Press). In sketching their "modular" view of the structure of the human mind, Sperber and Wilson (1986:71) begin as follows:

"Following Fodor (1983), we see the mind as a variety of specialised systems, each with its own method of representation and computation. These systems are of two broad types. On the one hand there are the input systems, which process visual, auditory, linguistic and other perceptual information."

On the other hand, Sperber and Wilson (1986:71) assume,

"there are the central systems, which integrate information derived from the various input systems and from memory, and perform inferential tasks."

As regards the input systems, Sperber and Wilson (1986:71) specifically assume that

"each input system has its own method of representation and computation, and can process only information in the appropriate representational format."

Accordingly, as Sperber and Wilson (1986:71) go on to point out,

"[a]uditory perception can process only acoustic information, and the processes involved in auditory perception differ from those involved in olfactory perception, etc."¹

Note in particular that language, on this view, is an input system along with the respective input systems for the five human senses.

These views of Sperber and Wilson's reflect Fodor's conception of 'modularity'. As Carston (1988:42) explicates and summarises this conception, all the input systems "share a cluster of properties which are characteristic of modular structures". She proceeds to distinguish five such characteristic properties, only the first of which concerns us directly here: modular structures are "special purpose" or "domain specific". This property, Carston (1988:42) takes it, has at least three aspects (only the first of which Fodor himself "explicitly mentions"): (a) modular structures "are sensitive to representations of only some subset of the full set of environmental stimuli which the organism is capable of processing"; (b) in each instance "the vocabulary and format of representation" which a modular structure manipulates "are specific to the particular system" - this is why no input system shares the vocabulary and format of its representations either with other input systems or with the central systems; (c) modular structures "each have a proprietary database and set of computational principles directing their processes".²

We can now turn to reconsider the question which arose earlier, namely "How do the processes at work in utterance interpretation relate to human cognition in general?" With reference to Sperber and Wilson's modular view of the structure of the human mind, at least two kinds of answer to this question are conceivable. In terms of one kind of answer, the processes at work in utterance interpretation operate in utterance interpretation only. In other words, they are "domain specific" or, equivalently, "special purpose" processes. In terms of the other kind of answer, the processes at work in utterance interpretation operate in all types of human cognition. In other words, they are "general purpose" processes.

One of the major claims of relevance theory is that, in utterance interpretation, "the real work of understanding" involves no processes of the special purpose, or domain specific, sort. The real work of understanding takes as its input the output from the process of

¹ S&W (1986) do not concern themselves explicitly with the question, "What type of mechanism according to Fodor provides the sensory information which the various input systems receive and process?" The type of mechanism which Fodor (1983) assumes for this purpose is what he calls *sensory transducers*. For a useful thumb-nail characterisation of sensory transducers cf. Carston 1988:41-42.

² The other four properties which according to Carston (1988:42) are considered characteristic of modular structures may be indicated as follows: the operations of modular structures are rapid; the operations of modular structures are automatic; modular structures have limited access to data (this is the property which Fodor has labeled *informational encapsulation*); finally, modular structures are "largely genetically specified" and, in consequence, they "mature at their own preprogrammed rate".

linguistic decoding.³ What the theory then claims is that, apart from linguistic decoding, utterance interpretation involves only the general purpose processes of human cognition. In terms of Fodor's conception of 'modularity', the claim here is that the processes of utterance interpretation are carried out by the central systems of the mind, because those processes are general purpose.

So relevance theory, as proposed by Sperber and Wilson, is first and foremost a theory of the general purpose processes of human cognition. It is not specifically a theory of utterance interpretation. Rather, the explanation of utterance interpretation in terms of relevance theory should be seen as a specific application of this theory.⁴ In the present chapter the account given of relevance theory will focus on the processing of information in general. The application of the theory to utterance interpretation will be taken up in chapter 4.

3.3 Relevance and information processing

3.3.1 The general goal of human cognition: efficient information processing

According to Sperber and Wilson (1986:47, 75, 150), the general goal of human cognition is to improve the individual's knowledge of the world, including his knowledge of himself. Such improvement takes at least three forms. Improvement of our knowledge of the world may come about, firstly, by the addition to it of more information. Improvement may come about, secondly, by an increase in the accuracy of our existing information. And, thirdly, improvement may come about by the integration of items of existing information into a more coherent whole so that each of the items will in future be more easily accessible when needed.⁵ In sum: the individual can improve his knowledge of the world by improving its *quantity*, its *quality* or its *organisation* (Sperber and Wilson 1987:700).

Sperber and Wilson (1986:49) claim that the human mind aims at the most efficient information processing possible. If so, then how is its general goal of threefold improvement of its knowledge of the world achieved as efficiently as possible? A short answer to this question can be given as follows. Firstly, cognition, including information processing, is relevance-oriented. Secondly, human beings accordingly always allocate their processing resources to the information which is most relevant: that information which is likely to bring about the greatest improvement in their knowledge of the world at the smallest processing cost. In other words, human beings aim at maximising the relevance of the information processed. They do this because that is the way the human mind works. Thirdly, the individual's specific cognitive goal at any given time is always consistent with the mind's more general goal of maximising the relevance of the information it processes.

³ Linguistic decoding, as noted above, is the task of an input system. S&W's view of the relationship between linguistic decoding and "the real work of understanding" in utterance interpretation is presented in more detail in 4.2 below.

⁴ For a discussion of reasons for viewing pragmatics as not itself constituting an input module cf. W&S 1986:67-71.

⁵ For some information on the notion of 'accessibility' cf. 4.3.1.

Inevitably, the short answer just given has been given in general and abstract terms. Let us now try to make this answer more specific and concrete.

3.3.2 Relevant information

As a rule, human beings cannot possibly process all the potential stimuli in their environment. Consider, for instance, the case of an individual who has woken up, and is standing at the window of his bedroom. The environment contains a myriad things to see, hear and smell; more formally speaking, it contains a vast range of visual, auditory and olfactory stimuli. From our own experience, we all know that in such circumstances one does not notice everything that there is to notice. That is, not all of the potential stimuli reach the level of one's attention. What determines which of the potential stimuli in an individual's environment do come to his attention? The answer given by relevance theory is: those stimuli from which the individual is most likely to gain relevant information. This answer leads, in turn, to the question: what makes information relevant?

In developing their answer to this question, Sperber and Wilson (1986:38) propose the notion of 'cognitive environment'. To this end, they first distinguish between a broad and a narrower sense of 'physical environment'. We human beings share a common physical environment in the broad sense in that "All humans live in the same physical world". Our physical environments in the narrower sense, however, may be very different from one person to another because of, for instance, geographical differences (as between, say, Iceland and Central Africa) or differences in workplace (as between, say, a cattle farm in the outback of Australia and a stockbroker's office in New York). But even if we human beings all shared the same narrower physical environment, Sperber and Wilson (1986:38) contend, as individuals we would still have different "cognitive environments". The differences in 'cognitive environment' between individual human beings are due, Sperber and Wilson (1986:38) assert, to factors such as the following: we human beings speak different languages, we have mastered different concepts, both our perceptual and our inferential abilities differ in effectiveness from one individual to another, and individuals have different memories.

Accordingly they propose the following notions of 'cognitive environment' and 'manifestness (of an assumption)' (Sperber and Wilson 1986:39-46):⁶

- (1) A *cognitive environment* of an individual is a set of assumptions which are manifest to him.
- (2) An assumption is *manifest* to an individual at a given time if and only if he is capable at that time of constructing it and accepting it as true or probably true.⁷

⁶ As was pointed out in n.8 to ch.2 above, assumptions, for S&W, are, informally speaking, "conceptual representations treated by the individual as representations of the actual world (as opposed to fictions, desires or representations of representations)" (S&W 1986:2).

⁷ For particulars of S&W's treatment of the notion of 'manifestness' cf. e.g. S&W 1986:39-46 and 196-200.

These definitions may be combined as follows: a cognitive environment of an individual at a given time is a set of assumptions (about the actual world) which the individual at that time is capable of constructing and accepting as true or probably true. In answer then to the question "What makes information relevant?" Wilson and Sperber (1986:72) claim that new information is relevant to someone if it interacts with his cognitive environment in any one of three specified ways. Relevance, in other words, is a relation between assumptions which an individual has recovered by processing some stimulus and assumptions which are contained in his cognitive environment. Let us consider the three ways in which new information can interact with an individual's cognitive environment, that is, with his existing assumptions about the world.⁸

Case 3.A

Suppose that you wake up with the following thought:

- (3) If it's raining, I'll stay at home.

To find out whether it is raining, you get out of bed, walk over to the window, and look outside. Of all the many visual stimuli potentially available for you to process, the sight of drops of water falling from the clouds to the ground is the one you are most likely to process at that point. This is so, because by processing that particular stimulus you obtain information which is highly relevant to you, namely:

- (4) It's raining.

This new information can combine with the "old" assumption (3) to yield further information, namely:

- (5) I'll stay at home.

This further information cannot be inferred from either the "old" assumption (3) alone or the new information alone. Rather, to infer this further information, you have to use the old assumption and the new information as joint premisses. What makes the new information relevant to you in the situation sketched above is precisely that it can conjoin with an existing assumption of yours to make such an inference possible. To be specific, the new information (4) can combine with the existing assumption (3) as the premisses of the following deductive argument:⁹

⁸ Presented here and in 3.3.3 is a slightly expanded version of the illustration provided by W&S (1986:72-74). For a rather more technical account of their notion of 'contextual effects' cf. S&W 1986:108-117.

⁹ We intend *deductive argument* to be taken in approximately the conventional sense presented in standard introductory textbooks of logic such as (Copi 1973) (or (Alexander 1969), where deductive arguments are referred to as *conclusive arguments*). To flesh out this remark in at least a minimal way, we will replace the concept of 'proposition' in (Copi's 1973:7) formulations by Sperber & Wilson's notion of 'assumption' (on the latter cf. n.8 to ch.2 above). We can then indicate the approximate senses in which we will be taking a number of elementary terms of logic. Firstly, we take an argument to be any group of assumptions "of which one is claimed to follow from the others, which are regarded as providing evidence for the truth of that one" (cf. Copi 1973:7). Secondly, we

- (6) If it's raining, I'll stay at home (= (3)).
It's raining (= (4)).
Therefore, I'll stay at home (= (5)).¹⁰

In relevance theory, an existing assumption such as (3) with which new information interacts in ways such as that just illustrated is called a *context*. New information such as (4), it is then said, *contextually implies* further information such as (5) in the context concerned. And inferred information such as (5) is said to be a *contextual implication* of the new information in the context concerned. New information is relevant, then, in any context in which it has contextual implications. Of course, too, the more contextual implications the new information has, the more relevant it is.

Note that in Case 3.A your knowledge of the world has been improved by the addition to it of more information.

Case 3.B

Suppose now that you wake up, again with the thought of (3). One of the potential stimuli for you to process is a pattering sound on the roof. By processing this auditory stimulus you derive new information which is highly relevant in the context of (3), namely the information in (7).

- (7) It's raining.

Again you walk over to the window. Again, among all the many potential visual stimuli, there is the sight of drops of water falling from the clouds to the ground. This stimulus will again affect you, the individual concerned, because it too yields new information which is relevant. In Case 3.B, however, its relevance is of a different nature from its relevance in

take the conclusion of an argument to be that assumption "which is affirmed on the basis of the other assumptions of the argument" (cf. Copi 1973:7). Thirdly, we take these other assumptions "which are affirmed as providing evidence or reasons for accepting the conclusion" to be the premisses of that argument (cf. Copi 1973:7). Fourthly, an argument is deductive (or, equivalently, demonstrative) if it "involves the claim that its premisses provide *conclusive* evidence" for the truth of its conclusion (cf. Copi 1973:20). Fifthly, every deductive argument is either valid or invalid: a deductive argument is valid "when its premisses, if true, do provide conclusive evidence for its conclusion, that is, when premisses and conclusions are so related that it is absolutely impossible for the premisses to be true unless the conclusion is true also" (cf. Copi 1973:20-21). Or as Black (1952:22) puts it: an argument is valid if "it is impossible for the premisses to be true without the conclusion also being true". In these terms, then, the deductive argument (6) above is valid. Sixthly, an argument is non-deductive/inductive/non-demonstrative/non-conclusive if it "involves the claim, not that its premisses give conclusive evidence for the truth of its conclusion, but only that they provide *some* evidence for it" (cf. Copi 1973:21). Sperber & Wilson (1986:passim, e.g. 37, 40-41, 81-117) take explicit account of these and related concepts of logic; they (1986:65-71) suggest, however, that "non-demonstrative inference, as spontaneously performed by humans, might be less a logical process than a form of suitably constrained guesswork".

¹⁰ Adapting some traditional logical terms, we will use the phrase a *minor premiss* to refer to an assumption/statement such as *It's raining* - expressed by means of a syntactically simple declarative sentence - and the phrase a *major premiss* to refer to an assumption/statement such as *If it's raining, I'll stay at home* - standardly expressed by means of a sentence of the *if, then* form (cf. e.g. Black 1942:133; Copi 1973:153-154).

Case 3.A. At the time when you look out of the window in Case 3.B, your cognitive environment already contains the assumption that it is raining: an assumption which you derived by processing an auditory stimulus (the pattering on the roof). But, because there are other possible explanations for the pattering sound, you cannot have complete confidence in the correctness of assumption (7). So, in this case the new information which you derive from the visual stimulus of raindrops falling is relevant because it raises your confidence in (7). To put it differently, the new information is relevant because it strengthens an existing assumption.¹¹ New information is relevant, then, in any context in which it strengthens one or more existing assumptions. The more assumptions it strengthens, and/or the more it strengthens these assumptions, the more relevant the new information is.

Case 3.C

Again, suppose that you wake up with the thought of (3), that you process the auditory stimulus of a pattering on the roof, and that you obtain by so doing the information that it is raining. So your cognitive environment now contains the assumption (4), *It's raining*. You walk to the window, and discover that although drops of water are falling on the roof, they are not coming from clouds in the sky. Instead, they are coming from a sprinkler on the lawn in front of the house. Processing the visual stimulus provided by the sprinkler on the lawn, and also the visual stimulus provided by the absence of rain clouds, you obtain the information in (8):

(8) It is not raining.

Intuitively, this new information too is highly relevant in the given context. Actually, however, this new information contradicts an existing assumption of yours. What is more, the new information conclusively shows this existing assumption to be mistaken.¹² In such a case, the new assumption is stronger than the old assumption. Accordingly, the old assumption has to be abandoned. New information is relevant, then, in any context in which it contradicts and disproves an existing assumption and so leads to its elimination. The more assumptions it contradicts and thereby eliminates, and/or the stronger the assumptions so eliminated, the more relevant the new information is.

Cases 3.A, 3.B and 3.C each illustrate one of the three ways in which new information can interact with, and so be relevant in, a context, where the context is a set of assumptions about the world held by an individual: (i) the new information can combine with existing assumptions to inferentially yield new assumptions; (ii) the new information can strengthen

¹¹ The underlying idea here is that factual assumptions - such as (7) - are entertained with greater or lesser confidence. That is, people think of them as more likely or less likely to be true. New information may then serve to improve our representations of the world by lowering or raising our confidence in some of our assumptions. The strength of an assumption, accordingly, depends on our degree of confidence in the assumption. For an extensive discussion of the notion 'strength of assumptions', and of how differing strength of assumptions affects the relevance of any new information, cf. S&W 1986:75-83, 108-117.

¹² As was mentioned in n.2 of ch.2, S&W interpret the term *information* in a broad sense, including false assumptions. In Case 3.C, the assumption (4), *It's raining*, is an example of a false assumption as an item of information in S&W's broad sense of the term.

existing assumptions; (iii) the new information can contradict and eliminate existing assumptions. The terms which Sperber and Wilson (1986:112, 115) use both in referring to these types of interaction, and in referring to the respective logical relations involved, are as follows: (i) *contextual implication(s)*, (ii) *(contextual) strengthening(s)*, and (iii) *(contextual) contradiction(s)*. Moreover the terms *contextual implication(s)* and *(contextual) strengthening(s)* are used in a third way as well: in referring to the results of the respective interactions. In the case of (contextual) contradiction(s), however, Sperber and Wilson (1986:115) refer to the associated result as *the erasure of premisses from the context*. And their (1986:115) umbrella term for referring to all three types of interaction grouped together is *contextual effects*. Accordingly, the cognitive effects of any new information are always contextual effects in this sense. And, in general, new information is relevant in any context in which it has contextual effects.¹³

Note that the notion of 'relevance' illustrated by Cases 3.A to 3.C is a comparative notion: the more contextual effects some new information has, the more relevant it is. This of course gives rise to at least one question: if relevance is a matter of degree, what is to prevent a representation produced by some input system from being endlessly processed in an ever-increasing context in order to yield more and more contextual effects? The answer given by relevance theory to this question is that relevance depends not only on contextual effects, but also on processing effort.

3.3.3 Processing effort

To see how processing effort can affect relevance, consider another illustration along the lines of the three presented above.

Case 3.D

Once again, you wake up and you think:

(9) If it's raining, I'll stay at home.

Suppose now, however, that *either* you walk to the window and notice:

(10) It's raining.

¹³ In Sperber & Wilson's (1986:260, n.25) view, information may be new either in the sense of "being new to the organism" or in the sense of "being newly processed": "In general, new information is not necessarily information that is new to the organism, but merely information that is being newly processed. Information retrieved from memory could be new information in this broader sense." In Cases 3.A to 3.D, the new information is information which "is new to the organism". Instances of information which is new in the sense of "being newly processed" are provided by, for example, the contextual assumptions (18) and (19) of 4.3.2 below: (18), *Peter is batting poorly*, is based in part on the hearer's memories of recent visual input and in part on items retrieved from the hearer's encyclopaedic knowledge of cricket; (19), *If a batsman is using a bat which is too big for him, this may cause him to bat poorly*, has in its entirety been retrieved from the hearer's encyclopaedic knowledge of cricket. (On Sperber & Wilson's notion of 'encyclopaedic knowledge' cf. 4.3.2 below, including in particular n.25.)

or you walk to the window and notice:

(11) It's raining and there are leaves on the lawn.

The new information (10) is obviously less extensive than the new information (11). Yet, intuitively, (10) is more relevant to you in the context of (9) than is (11). That is to say, the new information that there are leaves on the lawn has nothing to do with your choice, whether or not to stay at home, if this choice depends on the rain. But the intuition about (10)'s greater relevance cannot be explained on the basis of contextual effects: (10) and (11) have exactly the same contextual effects in this context. On the one hand, that is, both of them have the contextual implication (12).

(12) I'll stay at home.

On the other hand, (11) has no additional contextual effect that (10) does not have. The crucial difference is that the processing of (11) to derive this contextual effect requires *more effort* than does that of (10): since (11) contains (10) as a subpart, (11) requires *all* the processing effort required to process (10) *plus* some extra effort to process the other part of (11). And this extra effort yields no additional effect. We can say, then, that the extra effort required to process the new information (11) detracts from (11)'s relevance. Since in the context (9) the new information (10) yields exactly the same contextual effect as the new information (11), but does so in return for less processing effort, (10) is more relevant than (11).

3.3.4 A comparative definition of relevance

The illustration presented above as Cases 3.A to 3.D highlights two important points about Sperber and Wilson's notion of relevance. First, relevance is affected by a pair of factors -- namely, *the contextual effects produced* and *the processing effort required*. Second, relevance is a matter of degree: it depends on the *amount* of contextual effects produced, and on the *amount* of processing effort required. These two aspects of relevance are captured in the following comparative definition of relevance:¹⁴

(13) *Relevance*

- (a) Other things being equal, the greater the contextual effects produced, the greater the relevance achieved.
- (b) Other things being equal, the smaller the processing effort required, the greater the relevance achieved.

Let us return now to the claim, set out in 3.3.1, that the general goal of human cognition is to maximise relevance. This claim can now be made more precise with reference to the

¹⁴ This definition is adapted from W&S's (1986:74) definition. Both the latter definition and our adaptation of it are simplified versions of the technical definition presented in (S&W 1986:153). An important difference between these simplified versions and the technical version is that the latter defines relevance with reference to an individual. The significance of the notion that relevance is relevance for an individual will be taken up in 4.3.1.

definition of relevance presented above: human cognition is aimed at getting the greatest cognitive effects possible in exchange for as little processing effort as possible. As the illustration above has to some extent made clear, this means that we human beings must pay attention to those phenomena which, when represented in the best possible way and processed in the best possible context, seem likely to yield the greatest possible contextual effects in return for the processing effort available to us. But this immediately gives rise to the following question: how does any given individual "know" beforehand which phenomena are most likely to be worth his while to process? In the next section we consider a partial answer to this question.

3.4 The principle of relevance

Of course, as was noted in 3.3.2. above, the (cognitive) processing resources of every human being are limited. Consequently, the individual cannot process each and every potential stimulus in each and every possible context in order to determine its relevance. Such a wholesale, indiscriminate, approach to processing would consume large amounts of processing effort, and so would be at odds with the general goal of maximising relevance (specifically, it would fail to meet the requirement of minimal processing effort). Some sort of selection process, therefore, must be at work in the individual's mind to decide which phenomena it will process.

Ordinary phenomena carry no guarantee of relevance. That is, in attending to an ordinary phenomenon, the individual has no basis for expecting that processing it will yield sufficient cognitive effects to offset the processing effort required. There are certain "heuristics", however, which guide individuals in "picking out" phenomena likely to be relevant. As we all know from our own experience, certain phenomena seem automatically to pre-empt our attention. Consider for instance our reaction to a raging fire, to the crying of a baby, to a sudden loud noise, or to a smell of gas. These are examples of phenomena which are automatically processed by our perceptual systems. No doubt, human survival often depends on rapid responses to such phenomena. As Sperber and Wilson (1986:151-2; 1987:703) point out, some of the heuristics at work in selecting certain ordinary phenomena for automatic processing are innate - that is, an outcome of natural selection - and others are developed through experience. Presumably, for example, the way in which people automatically pay attention to all sudden loud noises is due to an innate heuristic, whereas the way in which parents' attention is pre-empted by the crying of a particular baby, even if barely audible, reflects a heuristic which has resulted from some form of learning (Sperber and Wilson 1986:151-152).

Of special interest here, however, are a certain class of stimuli that do come with a guarantee of relevance. The many potential stimuli in our environment include some which intentional organisms produce with the intention to inform their audience of something by having their audience recognise this intention to inform. Such stimuli Sperber and Wilson call *ostensive stimuli*. The ostensive stimuli produced by human beings include pointing at something with one's finger or arm, waving one or both hands, ringing a doorbell. Such stimuli are the means by which we as intentional organisms demand the attention of an

audience. Specifically, such stimuli focus the audience's attention on the communicator's communicative intention.¹⁵ But since humans pay attention to relevant information only, any attempt to get an audience's attention by means of an ostensive stimulus comes with a guarantee that the information which the producer of the ostensive stimulus intends to convey is relevant to the audience. In sum: when processing any ostensive stimulus directed to them, the audience are entitled to a firm expectation that they will find the stimulus relevant, because all such stimuli come with a tacit guarantee of their own relevance.

So far, these points about ostensive stimuli have been presented in mostly abstract terms; let us see if we can make them more concrete.

Case 3.E

Suppose you and a friend are walking down the street, when suddenly she stops and points at something over on the other side. By pointing, your friend demands your attention. Not only does she demand your attention; she also tries to direct or steer it. Specifically, she tries to focus your attention on her intention to inform you of something. But since humans pay attention to relevant information only, your friend, by intentionally claiming your attention, is giving a tacit guarantee that the information which she wishes to convey to you is relevant to you.¹⁶ Suppose that, in the direction in which your friend has pointed, you observe your dog trotting down the street. Suppose further that you would have expected your dog to be inside your house, and that your friend knows this. Then the information you derive by processing her ostensive stimulus is indeed sufficiently relevant to offset your processing effort. You know that, *generally speaking*, unsupervised dogs are liable to get run over in the streets, or to get themselves impounded. In the context of this knowledge, the new information that *your* dog is running around in the streets unsupervised will have clear contextual effects for you: clear contextual implications, for instance, that *your* dog at this very moment is in danger of being run over or impounded, and that *you* will have to act instantly to bring the animal back to safety. The information conveyed by the ostensive stimulus - in this case, your friend's pointing - is made relevant by these contextual effects.

The idea that an ostensive stimulus comes with a tacit guarantee of its own relevance - that is, with a tacit guarantee that processing it will be worth the audience's while - can be highlighted by means of an alternative scenario to the one just considered.

Case 3.F

Again, imagine that you and your friend are walking down the street. Suppose now that your friend stops, and that she stretches out her right arm, raises her right hand to about face-height, and straightens her right-hand index finger in the direction in which she is

¹⁵ S&W (1986:29) informally define a communicative intention as an intention to inform the audience of one's informative intention. The latter is the intention to inform one's audience of something. For a technical account of the notions of 'informative intention' and 'communicative intention' cf. S&W 1986:46-64.

¹⁶ For this use of the phrase *a tacit guarantee* cf. S&W 1986:49.

looking. Naturally enough, you take her to be making a conventional pointing-out gesture; so you stop too to have a look at whatever it is she must be pointing out. But suppose that, among all the many things visible to you, not one strikes you as especially significant. This is exactly what the situation will be, for instance, if the two of you have stopped at a spot well-known to you both, yet nowhere in all of it is there anything novel to be observed. Feeling puzzled, you next ask your friend explicitly, "What are you pointing at?"; she answers, "Nothing." It is likely that you will be taken aback by such a response - understandably so, because the ostensive stimulus in this case will have been presented to you in bad faith. Popularly speaking, the pointing-out will have been a fake.

Cases 3.E and 3.F illustrate, then, that ostensive stimuli have a "special status" in relation to the other, ordinary, stimuli in our environment. Ostensive stimuli automatically convey a tacit guarantee that the audience will find them relevant.¹⁷ Whereas one can only hope for relevance when attending to ordinary phenomena, one can have definite expectations of relevance when attending to ostensive stimuli.

One of the many questions that arise at this point is: what exactly is the tacit guarantee of relevance communicated by ostensive stimuli? As was noted in 3.3.1 above, Sperber and Wilson see it as the general goal of human cognition to maximise relevance - that is, to obtain maximal cognitive effects in return for minimal processing effort. Ostensive stimuli directed to an audience do not, however, carry a tacit guarantee of maximal relevance to the audience. Rather, they carry a tacit guarantee of *optimal* relevance to the audience. In essence, this means that an ostensive stimulus carries a tacit guarantee of *adequate* contextual effects for the audience in return for minimal processing effort by the audience. This guarantee tacitly guarantees, in other words, that the minimum of justifiable processing effort by the audience will be enough to make the ostensive stimulus yield a range of contextual effects which will be adequate from the audience's point of view. Or to put it more succinctly: every ostensive stimulus carries a tacit guarantee that the audience will find it optimally relevant. In the principle of relevance, which Sperber and Wilson (1986:158) formulate as follows, this tacit guarantee is implicitly provided for by means of the term *presumption*:

(14) *Principle of relevance*

"Every act of ostensive communication communicates the presumption of its own optimal relevance."

The principle of relevance, it should be noted, does not claim that every ostensive stimulus actually fulfils the tacit guarantee of its optimal relevance. As was illustrated above by the scenario of the "fake" pointing-out, a communicator using an ostensive stimulus can give a tacit guarantee of optimal relevance in bad faith. Or she may give it mistakenly, for instance by presenting certain information to an audience in the mistaken belief that the audience do

¹⁷ For this use of the phrase *automatically convey* cf. W&S 1989:109.

not already have it. As Wilson and Sperber (1989:109) point out quite emphatically:

"The principle of relevance ... is not something that people ... obey or might disobey: it is an exceptionless generalization about human communicative behaviour ...

"[Communicators - M.S.-W.K.W.] may try hard or not at all to be relevant to their audience; they may succeed or fail; they still convey a presumption of relevance: that is, they convey that they have done what is necessary to produce an adequately relevant [ostensive stimulus - M.S.-W.K.W.]." ¹⁸

So, what the principle of relevance does claim is that there is no way for a communicator to avoid being taken by her audience as communicating a tacit guarantee that the ostensive stimulus which she is using is going to be found by the audience to be optimally relevant.

We said above that in the wording of the principle of relevance the tacit guarantee of optimal relevance is provided for implicitly by means of the term *presumption*. Let us see now what content Sperber and Wilson give to the presumption of optimal relevance.

3.5 The presumption of optimal relevance

The presumption of the optimal relevance of the ostensive stimulus used has two parts: one relating to the effect side of relevance, and the other to the effort side.

On the effect side, the presumption is that the ostensive stimulus, when processed by the audience in a context which the communicator expects them to have accessible, will yield enough contextual effects to make the ostensive stimulus worth the audience's attention. What is presumed, in other words, is a guarantee of the adequacy of the contextual effects. But how much is *enough* or *adequate*? Sperber and Wilson (1986:160-1) claim that this factor is variable. What counts as adequate cognitive effects will vary from audience to audience, and from occasion to occasion. Sperber and Wilson suggest, for instance, that for various social occasions the conventionally expected level of relevance will be differently defined within a given culture.

On the effort side, the presumption is that the communicator has chosen an ostensive stimulus which will not put her audience to any unjustifiable processing effort in their recovery of her intended contextual effects. In other words, the communicator must choose the ostensive stimulus that will make her intended contextual effects as easy as possible for her audience to recover. ¹⁹

¹⁸ These points are likewise emphasised in (S&W 1986:162). S&W (1986:50) also characterise the principle of relevance as being a "fact" - namely, the fact "that an act of ostension carries a guarantee of relevance".

¹⁹ In regard to the effort required from the communicator herself, S&W (1986:157) "assume that she eliminates any stimuli which would require too much effort on her [own - M.S.-W.K.W.] part (e.g. drawing a map when a verbal indication will do)..." According to S&W (1986:157), the choice of a stimulus by a communicator is also constrained by considerations not related to processing effort. An

Accordingly the presumption of optimal relevance, which is taken by the audience/addressee as being automatically communicated by every ostensive stimulus, is worded by Sperber and Wilson (1986:158) as follows, where "addressee" is interchangeable with "audience", and where "the set of assumptions {I}" denotes the "something" which the communicator intends to convey to the addressee:

(15) *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}."

For ease of reference, we introduce here two terms of our own as labels for the two parts of the presumption of optimal relevance: it will be convenient to be able to refer to part (a) as *the adequate effect presumption* and to (b) as *the least effort presumption*.

3.6 Consistency with the principle of relevance

So far in chapter 3 we have focused on relevance theory as a theory of non-verbal ostensive communication - that is, communication by means of non-verbal ostensive stimuli.²⁰ The heart of the theory, as we have seen, is the principle of relevance; this is the "exceptionless generalisation"²¹ that "every act of ostensive communication communicates the presumption of its own optimal relevance".

But our main concern is of course with relevance theory as it applies to the description and the explanation of verbal communication. Of crucial importance to us, therefore, is the fact that all linguistic utterances are ostensive stimuli. That is, every utterance is produced by a communicator who is an intentional organism, and who produces the utterance with the intention to inform the audience of something. The linguistic communicator moreover intends, by means of the utterance, to make her audience recognise her intention to inform. Everything said above about ostensive stimuli applies therefore to linguistic utterances as well.

Accordingly, let us now review some more of the main ideas of relevance theory presented thus far but, in doing so, let's apply them to linguistic utterances.²² Firstly, then, in regard to (linguistic) utterances the principle of relevance implies the following: every utterance

example of this is provided by cultural rules which prohibit the use of certain words; the communicator may well avoid such words as being objectionable.

²⁰ Pointing at something with one's finger or arm, waving one or both hands, and ringing a doorbell were mentioned in 3.4. as instances of non-verbal (or, equivalently, non-linguistic) ostensive stimuli. Such stimuli are of a non-verbal nature in that, to produce them, the communicator in principle need not and in practice generally does not speak or write some language unit such as a sentence, phrase, or word.

²¹ This characterisation of the principle of relevance was first quoted in 3.4 above.

²² We here draw heavily, and gratefully, on an outline given by Smith (1989:172).

addressed to a hearer conveys to the hearer automatically a tacit guarantee that he is going to find the utterance optimally relevant to himself.²³ Secondly, the utterance is relevant if and only if it has contextual effects - that is, if and only if it allows the hearer to infer conclusions which would follow neither from the utterance alone nor from the context alone.²⁴ Thirdly, the context is a set of assumptions about the world which the hearer holds (and with which the content of the utterance interacts in one of three specified ways).²⁵ Fourthly, the utterance is optimally relevant if and only if it enables the hearer to achieve an adequate range of contextual effects, but puts the hearer to no unjustifiable processing effort in achieving these contextual effects.²⁶

Fifthly, the fact that an utterance conveys to the hearer automatically a tacit guarantee of its being optimally relevant to him does not imply that it will actually be optimally relevant to the hearer.²⁷ On the one hand, for instance, the speaker may tell the hearer something in the mistaken belief that he does not already know it. On the other hand, the tacit guarantee may be given in bad faith. That is, the speaker may be speaking merely in order to distract the hearer's attention from relevant information elsewhere. In either of these two situations the hearer will be unable to find an interpretation which justifies the guarantee (Wilson and Sperber 1987a:14). In short: despite the tacit guarantee of optimal relevance, an utterance may fail to communicate. In this connection, Sperber and Wilson introduce the concept of '(an utterance's being) consistent with the principle of relevance': an utterance on a given interpretation is consistent with the principle of relevance if and only if a rational speaker might have expected/foreseen that the utterance on that interpretation would furnish the hearer with an adequate range of contextual effects in return for the minimum of justifiable processing effort.^{28,29}

Consider again the infinitival phrase used in the last sentence of the previous paragraph, *to furnish the hearer with an adequate range of contextual effects in return for the minimum of justifiable processing effort*. The act of expecting/foreseeing attributed to a rational speaker by means of this phrase is, of course, the act of expecting/foreseeing the utterance's being optimally relevant to the hearer. So, if this sense of the *to*-phrase is expressed in terms of Sperber and Wilson's technical concept of 'optimally relevant'/'optimal relevance', their further technical concept of 'consistent/consistency with the principle of relevance' can be

²³ For a statement of what the principle of relevance claims in regard to ostensive stimuli in general cf. 3.4 above.

²⁴ For an exposition of this main idea of relevance theory cf. 3.3 above.

²⁵ This notion of 'context' was presented in connection with Case 3.A in 3.3.2 above.

²⁶ This notion of 'optimal relevance' was presented in 3.5 above.

²⁷ In presenting this fifth main idea of relevance theory, we closely follow the exposition given by W&S (1987a:14).

²⁸ Cf. S&W 1986:166; W&S 1986:76; W&S 1987a:14, and also Smith 1989:172. We give *expect* and *foresee* as synonymous alternatives because W&S (1986:67-77; 1987a:14-16) seem to us to use these verbs interchangeably in combinations such as *a rational communicator might have expected*, and *a rational communicator might have foreseen*.

²⁹ So far as we know, S&W have not published an explicit and systematic characterisation of the notion of 'rationality' which they take to be operative here.

more briefly defined as follows:

(16) *Consistency with the principle of relevance*

An utterance on a given interpretation is consistent with the principle of relevance if and only if a rational speaker might have expected/foreseen that the utterance on that interpretation would be optimally relevant to the hearer.³⁰

Sperber and Wilson make a strong and startling claim to which the above-defined property is central.³¹ This claim may be stated as follows:

(17) Every linguistic utterance has at most one interpretation which is consistent with the principle of relevance.

The import of this claim will be illustrated in the course of the next chapter, chapter 4. For the present, let us bring up two points, however. First point: this claim implies that some linguistic utterances have no interpretation - *none at all* - which is consistent with the principle of relevance. Instances illustrating this are given in connection with, for example, Cases 4.F and 4.G below; in such instances, verbal communication fails. Second point: this claim asserts that a linguistic utterance which does succeed in communicating will never have *more than one* interpretation which is consistent with the principle of relevance.³²

In chapter 4 we will review in some detail how, according to Sperber and Wilson, the principle of relevance serves as the basis for an explanatory account of the various aspects of utterance interpretation outlined in chapter 2.

³⁰ With reference to ostensive communication in general, here is how S&W (1986:166) define their 'consistency' concept: "Let us say that an interpretation is *consistent* with the principle of relevance if and only if a rational communicator might have expected it to be optimally relevant to the addressee."

³¹ For this claim cf. e.g. W&S 1986:76 and 1987a:14. In regard to ostensive communication in general W&S (1986:167) state this claim as follows: "the principle of relevance does not generally warrant the selection of more than one interpretation for a single ostensive stimulus".

³² Two of the key words used in stating these two points are *none* and *one*. So as a mnemonic by which to remember the import of this claim one might use the label *The "none or one" claim*.

Chapter 4

HOW THE CRITERION OF CONSISTENCY WITH THE PRINCIPLE OF RELEVANCE EXPLAINS UTTERANCE INTERPRETATION

4.1 Introduction

In this chapter we turn to a specific application of relevance theory, namely the explanation of utterance interpretation. The structure of chapter 4 "mirrors" the structure of chapter 2. For each paragraph in chapter 2 in which an aspect of the process of utterance interpretation is identified and illustrated, there is a corresponding paragraph in chapter 4 in which we explicate how, according to Sperber and Wilson, the principle of relevance can serve as the basis for an explanatory account of this aspect. In the last paragraph, 4.9, we conclude by outlining the common pattern for explaining the various aspects of utterance interpretation in terms of relevance theory.

4.2 The role of decoding in utterance interpretation

As we noted in 2.2 above, one of the major subtasks of utterance interpretation is to recover the semantic representation(s) associated with the underlying sentence. In the present paragraph, our aim is to set out how, in Sperber and Wilson's view, this subtask fits into their overall theory of utterance interpretation. To this end, we start off by considering, in rather more detail than we did in 2.2, the recovery of the semantic representations associated with the sentence underlying the utterance *Peter's bat is too big* (first presented as (4) in 2.3.2 above).

Consider the complex of English speech sounds roughly transcribed below:

(1) [pi:təzbætɪztvbig]

Consider too the task of interpretation to be carried out by the mind of a native speaker-hearer of English who is aware that another native speaker-hearer of English has addressed an utterance of this complex of speech sounds to him. (Switching back to our usual terminology, we will refer to these two people as *the speaker* and *the hearer*, respectively.) In appropriate conditions, the hearer's mind will automatically construct an assumption of the form in (2) (Sperber and Wilson 1986:177):

(2) The speaker has uttered the sentence *Peter's bat is too big*.

Because the utterance is expressed in the hearer's native language, his mind automatically analyses it as a token of a particular linguistic structure, here the sentence *Peter's bat is too big*, and automatically assigns to the utterance the semantic representations which the generative grammar of English links to the sentence uttered (Sperber and Wilson

1986:177).¹ The grammar links two semantic representations to this sentence, and these differ in a single respect only: one of them incorporates in the subject phrase the noun *bat* in the sense of (roughly) 'hitting instrument' and the other, the noun *bat* in the sense of (roughly) 'flying mammal'.² For the rest - namely, as regards the meaning of *Peter's* and the meaning of *is too big* - the two sentence meanings, and therefore the two semantic representations as well, are of course the same. In very broad terms, both sentence meanings involve the idea that some entity has some characteristic/property. Some information about the entity's identity is contained in *Peter's bat*, the subject phrase of the sentence. The predicate phrase of the sentence, *is too big*, contains some information about the nature and the extent of the property being attributed to the entity.

Firstly, then, what is the meaning of *Peter*? Our tentative answer to this question derives from remarks quoted earlier.³ Sperber and Wilson (1986:10), that is, remark that the generative grammar of a language cannot determine to which entities the proper names of the language refer, but can only provide some very general indications - for instance, by stating that a proper name such as *Peter* refers to a person or other entity bearing that name. As a result, a proper name such as *Peter* in *Peter's bat* is referentially indeterminate.⁴ Accordingly, we take it that the answer to the question above is as follows:

- (3) *Peter* means (approximately) 'some entity bearing the name *Peter*'.

Secondly, what is the meaning of the genitive '*s* in *Peter's bat*? Here, our answer reflects Sperber and Wilson's (1986:188) point that the meaning of the English genitive is essentially vague.⁵ It is because of this semantic vagueness that an utterance of the noun phrase *Peter's bat* can be used to refer, for instance, to a bat belonging to or being used by Peter⁶, or to a bat bought or chosen by Peter - to mention just four of the indefinitely many interpretations which could be cited as illustrating the vagueness of the meaning of the genitive '*s*. So here is our answer to the second question above:

- (4) The genitive '*s* means (approximately) 'somehow associated with'.

Using the meaning characterisations sketched in (3) and (4) for *Peter* and the genitive '*s*, respectively, we can now indicate the differing phrasal meanings which the hearer's mind is able to assign to an utterance of the noun phrase *Peter's bat* on the basis of the generative grammar of English. On the one hand there is the phrasal meaning which springs in part from the noun *bat* in the sense of 'hitting instrument':

- (5) *Peter's bat* can mean (approximately) 'some hitting instrument somehow associated with some entity bearing the name *Peter*'.

¹ The implicit assumption here is that the grammar of a language is somehow represented in the mind of the native speaker-hearer. So "grammar" here is short for "mental grammar".

² On these two senses and on the semantic ambiguity to which they give rise cf. 2.3.2 above.

³ These remarks were quoted in 2.3.3 above.

⁴ On the notion of 'referential indeterminacy' cf. 2.3.3 above.

⁵ On this point cf. 2.3.4 above.

⁶ More strictly speaking, *Peter* here should be replaced by something like 'some entity bearing the name *Peter*'. We have opted for the shorter, if less accurate, wording.

On the other hand there is the phrasal meaning which arises in part from the noun *bat* in the sense of 'flying mammal':

- (6) *Peter's bat* can mean (approximately) 'some flying mammal somehow associated with some entity bearing the name *Peter*'.

But, what, thirdly, is the meaning of the predicate phrase *is too big*? Because *is* is in the Present Tense, the time reference of this phrase is to the "now" of the particular utterance; that much seems uncontroversial. The adjective *big* denotes a property/characteristic. The adverb *too* denotes an excessive amount (of that property). So the sentence meanings both involve the idea that the bat concerned, whatever it is, has, at the time of utterance, some property of bigness to an excessive extent. But note again that *too* is semantically incomplete.⁷ If a bat is being (cryptically) said to be too big, of course the idea conveyed and grasped is that the bat is too big *for something*. The phrase *is too big*, however, contains no information about the nature of this reference-point. In short, then, our tentative answer to the third question goes as follows:

- (7) *is too big* means (approximately) 'has, at the time of utterance, some property of bigness to an extent which is excessive relative to some unidentified reference-point'.

As we noted at the start of this section, the hearer's mind automatically assigns two semantic representations to an utterance of *Peter's bat is too big*. The respective contents of these semantic representations can, we are assuming, be roughly sketched by combining, on the one hand, the meanings sketched in (5) and (7) above and by combining, on the other hand, the meanings sketched in (6) and (7).⁸ Here is the outcome of the first combination of meanings (drawing on (5) and (7)):

- (8) 'Some hitting instrument somehow associated with some entity bearing the name *Peter* has, at the time of utterance, some property of bigness to an extent which is excessive relative to some unidentified reference-point.'

And here is the outcome of the second combination of meanings (drawing on (6) and (7)):

- (9) 'Some flying mammal somehow associated with some entity bearing the name *Peter* has, at the time of utterance, some property of bigness to an extent which is excessive relative to some unidentified reference-point.'

What type of process is it that enables the hearer to assign semantic representations such as (8) and (9) to an utterance of the sentence *Peter's bat is too big*? Sperber and Wilson claim that this process is a *decoding* process. A decoding process, according to them (1986:13), "starts from a signal and results in the recovery of a message which is associated to the signal by an underlying code."

⁷ This point was first brought up in 2.3.4 above.

⁸ Recall that both (5) and (6) incorporate the meanings sketched in (3) and (4) for *Peter* and the genitive 's, respectively.

In the (automatic) process of linguistic decoding by the hearer's mind, the underlying code is the generative grammar of the particular language; what this grammar does is "pair phonetic and semantic representations of sentences" (Sperber and Wilson 1986:9). In the case of linguistic decoding moreover the "signal" is an utterance - for instance, an utterance of the complex of speech sounds roughly transcribed in (1) above - and the "message" is a semantic representation - for instance, (8) or (9) above - which the generative grammar of the language assigns to the particular linguistic structure uttered. (In our example above, the particular linguistic structure is of course the English sentence *Peter's bat is too big.*)

Sperber and Wilson emphasise the automatic nature of linguistic decoding. They (1986:177) hold that

"a linguistic stimulus triggers an automatic process of decoding ... we cannot choose to hear an utterance in a language we know as merely an unanalysed stream of sounds. We automatically recover its semantic representation, even if we accidentally overhear it and know it was not meant for us ... The linguistic decoding system has all the hallmarks of automatic, reflex perceptual systems such as hearing and vision. In the terms of Fodor ... it is an input system rather than a central processing system."⁹

The view that the semantic representations associated with an utterance are recovered by decoding implies, of course, that the contents of these semantic representations are determined solely by the generative grammar of the language in question. These contents, according to Sperber and Wilson (1986:175), are not affected by "the interests or point of view of the hearers". All native speaker-hearers of English, in all situations, will therefore recover the same two semantic representations - with the respective contents roughly sketched in (8) and (9) - for an utterance of the sentence *Peter's bat is too big.*¹⁰

What, then, is Sperber and Wilson's view on the role played by linguistic decoding in utterance interpretation or, equivalently, in utterance comprehension? Because Sperber and Wilson (1986:177) take utterance interpretation to be a process of identifying the speaker's informative intention, they hold the view that

"linguistic decoding is not so much part of the comprehension process as something that precedes the real work of understanding...[linguistic decoding - M.S.-W.K.W.] merely provides an input to the main part of the comprehension process."

At least two questions arise at this point. The first question: in what ways do the semantic representations recovered by linguistic decoding fall short of representing the speaker's

⁹ On the intended distinction between input systems and central processing systems cf. 3.2 above.

¹⁰ The fact that the use of natural language(s) is based (in part) on a code constitutes a major difference between the interpretation of linguistic stimuli and that of other types of ostensive stimuli. Briefly: because the use of linguistic stimuli is (in part) code-based, linguistic communication is more explicit and therefore capable of being more precise. S&W accordingly see linguistic communication as a "stronger" variety of ostensive communication than are the non-code-based varieties. For a discussion of this difference in "strength" cf. S&W 1986:174-175.

informative intention? The second question: what is the nature of the "main part of the comprehension process"?

The first question can be answered with reference to, for example, the semantic representations sketched in (8) and (9) above. Firstly, sentence meanings are sets of semantic representations, and for each sentence of the language this set contains as many semantic representations as the sentence has distinct senses/meanings (Sperber and Wilson 1986:193; cf. also 175, 177).¹¹ For example, the sentence *Peter's bat is too big* is ambiguous between the 'hitting instrument' and 'flying mammal' senses of *bat*. Accordingly the meaning of that sentence is a set of two semantic representations - the semantic representations sketched in (8) and (9) above.

Secondly, Sperber and Wilson (1986:193) emphasise that "[s]emantic representations are incomplete logical forms, i.e. at best fragmentary representations of thoughts". Semantic representations can be incomplete, in the sense intended here, in that they may contain indeterminate referring expressions such as not only proper names - for example, *Peter* in *Peter's bat is too big* - but also pronouns.¹² Semantic representations can also be incomplete in that they may contain "underdefined constituents" such as *too* and the genitive '*s*'.¹³

Thirdly, owing to referential indeterminacy such as that of *Peter*, and to semantic incompleteness/underdefinition such as that of *too*, a single sentence, in conjunction with a single one of its semantic representations, can express an unbounded range of thoughts.¹⁴ This has the effect that there is a gap between the meaning of a sentence (as represented in its semantic representations) and the intended interpretation of an utterance of that sentence (Sperber and Wilson 1986:10).

Fourthly, referential indeterminacy and semantic incompleteness are only two members of the large set of factors responsible for the gap by which the meaning of a sentence is separated from the intended interpretation of an utterance of that sentence. As was illustrated in 2.4, the intended interpretations of utterances can have a variety of aspects which are implicit - aspects, that is, which the semantic representations of the sentences uttered do *not* represent *at all*. For instance, neither of the semantic representations sketched in (8) and (9) contains any information on the speaker's attitude to the thought expressed in these representations, or on the intended illocutionary force of the utterance, or on any intended implicatures of the utterance.

¹¹ S&W (1986:73, 175, 193) seem to use *sense* and *meaning* interchangeably in speaking of "the sense of a sentence", "every sense of the sentence" and "the meanings of sentences".

¹² The referential indeterminacy of *Peter* as a referring term of English is reflected in (3) above. Examples of pronouns as indeterminate referring expressions of English are *she*, *her* and *it* in S&W's (1986:73) example sentence *She carried it in her hand*.

¹³ The vagueness, or underdefinition, of *too* as an adverb of English is reflected in (7) above by the words "relative to some unidentified reference-point" and that of the English genitive '*s*' in (4) above by the words "somehow associated with". The vagueness of *some time* as an underdefined constituent was illustrated in 2.3.4 above.

¹⁴ S&W 1986:10 make this point with reference to "a single sentence with a single semantic representation," but obviously the point extends to the case we mention in the text.

From what has just been pointed out, it is clear that, even in the case of a seemingly simple utterance such as *Peter's bat is too big*, the semantic representation recovered by linguistic decoding is far from representing the speaker's informative intention or, equivalently, the intended interpretation of the utterance. This brings us back to the second question posed above: what is the nature of "the main part of the comprehension process"? This question can now be slightly rephrased: what is the nature of the process which is capable of "bridging" the gap between the information contained in the semantic representations associated with an utterance and the intended interpretation of that utterance? It is, Sperber and Wilson (1986:176) claim, an *inferential process*. According to them (1986:12-13), "[an] inferential process starts from a set of premisses and results in a set of conclusions which follow logically from, or are at least warranted by, the premisses." Sperber and Wilson (1986:13) emphasise that, "[in] general, conclusions are not associated to their premisses by a code, and signals do not warrant the messages they convey."

When this inferential process which makes up the main part of the comprehension process operates correctly on the correct input from the linguistic decoding process, the result produced is the correct interpretation of the utterance concerned. Sperber and Wilson claim that relevance theory is able to explain this successful operation of the process of utterance interpretation. In general, their theory claims that "the correct interpretation of an ostensive stimulus is the first accessible interpretation consistent with the principle of relevance" (Sperber and Wilson 1986:177). (Recall that an interpretation is consistent with the principle of relevance if and only if a rational communicator might have expected/foreseen that that interpretation would be optimally relevant to the audience/addressee.¹⁵) The notion 'first accessible interpretation' can roughly be characterised as 'the first interpretation to occur to the hearer's mind'.¹⁶

Note that in this general characterisation of a correct interpretation, the property of consistency with the principle of relevance is made to function as the criterial property - indeed, *the sole and sufficient criterial property - by which to identify the correct interpretation of an ostensive stimulus*. In regard to the interpretation of linguistic utterances as a subclass of ostensive stimuli, Sperber and Wilson (1986:177) consequently put forward the more specific claim statable as below, where "first" is interchangeable with "most readily":

(10) *The consistency criterion*

The correct interpretation of a linguistic utterance is the first accessible interpretation of it which, on being appropriately tested, is found to be consistent with the principle of relevance.

They label this claim *the criterion of consistency with the principle of relevance*, but on the whole we find it convenient to use a shorter label - namely, *the consistency criterion*.

¹⁵ On S&W's concept of 'consistency with the principle of relevance' cf. 3.6 above.

¹⁶ The notion of 'accessibility' is in fact much more complex. For a somewhat more detailed discussion cf. 4.3.2 below.

So according to Sperber and Wilson the inferential task which the hearer of an utterance performs in order to go from the semantic representation(s) of the sentence uttered to the interpretation intended by the speaker is governed by the consistency criterion. This view of the nature of the hearer's inferential task in interpreting linguistic utterances will be extensively illustrated in the remainder of chapter 4.

4.3 How the consistency criterion explains the recovery of the explicit content of utterances

4.3.1 Introduction

In order to determine the intended interpretation of an utterance, a hearer must recover the intended explicit content of the utterance. Equivalently, the hearer's task - as we noted in 2.3 - can be said to include, as a major subtask, the recovery of the proposition expressed by the utterance. In these two sets of terms, then, what is recovered by linguistic decoding - as we saw in 4.2 - is neither a representation of the intended explicit content of the utterance, nor a representation of the proposition expressed by the utterance. Sperber and Wilson's technical label for such a representation is a *(unique) propositional form*. Let us see how this label is to be understood.¹⁷

To begin with, recall that by *thoughts* Sperber and Wilson (1986:2) mean "conceptual representations (as opposed to sensory representations or emotional states)", and that by *assumptions* they mean "thoughts treated by the individual as representations of the actual world (as opposed to fictions, desires, or representations of representations)."¹⁸ A logical form, for Sperber and Wilson (1986:72), is a conceptual representation from which all its non-logical properties - such as being happy or sad, or being "located in a certain brain at a certain time for a certain duration" - have been abstracted away. Such a logical form is "semantically complete", or "propositional", if it represents "a state of affairs, in a possible or actual world, whose existence would make it true" (Sperber and Wilson 1986:72). And if a logical form is semantically complete in this sense, it is "capable of being true or false". If a logical form is not propositional or, equivalently, is not semantically complete, it is "non-propositional". As a psychological example of a non-propositional, or semantically incomplete, logical form Sperber and Wilson (1986:72) mention the sense of a sentence. They (1986:72-73) illustrate this as follows:

"Given that 'she' and 'it' in (2) below do not correspond to definite concepts, but merely mark an unoccupied place where a concept might go, sentence (2) is neither true nor false:

(2) She carried it in her hand."

¹⁷ For this use of the notion 'unique propositional form' cf. S&W 1986:179.

¹⁸ Sperber and Wilson's (1986:2) informal definitions of the notions of 'thought' and 'assumption' were first mentioned in n.8 to ch.2 above.

Similarly, since both (8) and (9) above each represent (very roughly) a sense of a sentence - namely, of the sentence *Peter's bat is too big* - each of them is an incomplete logical form. To expand, say, (8) into the semantically complete, or propositional, (logical) form intended by the speaker of an utterance of *Peter's bat is too big*, the hearer at the very least has to take the following steps of interpretation: he must resolve the referential indeterminacy of *Peter* by pinpointing the specific individual referred to, he must enrich the vague meaning of the genitive 's sufficiently, he must decide between the 'hitting instrument' and 'flying mammal' senses for *bat*, he must fix the time reference of the Present Tense of the verb *is* in relation to the 'now' of the specific time of utterance, and he must identify the standard or reference-point relative to which *too big* is to be interpreted. Only if the hearer takes at least these steps of interpretation can he "bridge" the "gap" between the semantically incomplete logical form (8) and the unique propositional form intended by the speaker.¹⁹

So, for communication to succeed, the hearer must identify the correct propositional form - that is, the propositional form intended by the speaker. But by what criterion is the hearer to identify the correct propositional form? Here, Sperber and Wilson suggest an additional criterion of consistency, partly analogous to the consistency criterion (10). Recall that in terms of the consistency criterion (10) the hearer is to take as correct, for a linguistic utterance, the first accessible interpretation which is consistent with the principle of relevance. In terms of the additional criterion suggested by Sperber and Wilson (1986:184), the hearer is to take as correct, for a linguistic utterance, the propositional form which "leads to an overall interpretation which is consistent with the principle of relevance." The additional criterion, be it noted, operates in tandem with the consistency criterion (10). Indeed, its operation is dependent on (10): it does not declare a propositional form correct, unless that propositional form is (known to be) instrumental in producing an overall utterance interpretation which meets the consistency criterion (10). Furthermore, according to Sperber and Wilson (1986:184), if a propositional form qualifies as correct in terms of the additional criterion, it is "itself ... consistent with the principle of relevance".

But this gives rise to a further question: what *general procedure* might the hearer use to identify the propositional form that meets the additional criterion of consistency with the principle of relevance? As we turn now to consider Sperber and Wilson's answer to this question about a general procedure, let us keep in mind the point - noted in 2.3.1 and 2.5 above - that to identify a propositional form is to carry out three asks: disambiguation, reference assignment, and enrichment. Here, then, is how Sperber and Wilson (1986:185) state their answer to the above-posed question about a general procedure for the hearer:

"At every stage in disambiguation, reference assignment and enrichment, the hearer should choose the solution involving the least effort, and should abandon this solution only if it fails to lead to an interpretation consistent with the principle of relevance."

¹⁹ For *unique* in this connection cf. S&W 1986:179.

The less the effort required to process an interpretation, the more accessible the interpretation is, and conversely.²⁰ We therefore assume that Sperber and Wilson's answer quoted above is an informal way of saying the following: throughout each of the three tasks which he carries out in identifying the correct propositional form, the hearer should be guided by an additional criterion which "meshes" with, and for its operation is dependent on, the consistency criterion (10). That is, throughout the disambiguation task, the reference assignment task and the enrichment task, the hearer should invariably select the first accessible interpretation - or, equivalently, the most readily accessible interpretation. Note what this involves:

- * in relation to *each semantic ambiguity*, the hearer should select the most readily accessible interpretation which resolves the ambiguity;
- * in relation to *each indeterminate referring expression*, the hearer should select the most readily accessible interpretation in terms of which a referent is assigned to the expression;
- * in relation to *each sentence constituent whose meaning is vague or underdefined*, the hearer should select the most readily accessible interpretation which enables that constituent's meaning to be enriched.

We will refer to such interpretations as *task-specific interpretations*.

Now, suppose that a first accessible interpretation is selected as a task-specific interpretation and, as such, leads to an overall interpretation which meets the consistency criterion. Then, in terms of the additional criterion, the hearer must assume that this task-specific interpretation is the correct one. Only if the first accessible task-specific interpretation fails to lead to an overall interpretation which meets the consistency criterion should the hearer abandon the first accessible interpretation and proceed to consider a less accessible, and therefore more effort-consuming, alternative.

Here, then, is the picture that has emerged so far. The additional criterion governs the way in which the hearer identifies correct task-specific interpretations. These are interpretations of specific aspects of the overall interpretations of utterances. The additional criterion fulfils its function by operating in tandem with the consistency criterion (10). The additional criterion can be stated and labeled as follows, where again "first" is interchangeable with "most readily":

(11) *The task-specific criterion*

The correct task-specific interpretation is the first accessible task-specific interpretation which leads to an overall interpretation which meets the consistency criterion (10).

Of course (11), like (10), is a criterion which, for speaker and hearer alike, makes consistency with the principle of relevance the criterial property by which to recognise the correctness of interpretations. The two criteria apply, however, to interpretations of

²⁰ On this correlation between (on the one hand) the mental effort required to construct an interpretation and (on the other hand) the accessibility of that interpretation to the mind cf. 4.3.2 below.

different "sizes". Criterion (11) concerns the correctness of task-specific interpretations, which are entities of a lesser size than overall interpretations, with whose correctness criterion (10) is concerned. More elaborately speaking, one could accordingly label (10) "the consistency criterion for overall interpretations" and label (11) "the consistency criterion for task-specific interpretations". For the sake of relative brevity, however, we will in what follows refer to (11) as *the task-specific criterion*.

The task-specific criterion, be it also noted, extends the concept of 'consistency with the principle of relevance' and the concept of 'correct (as an interpretation)'. The existing senses of these two concepts may be stated as follows:

- (12) An overall interpretation of a linguistic utterance is consistent with the principle of relevance, and hence correct, if that interpretation might have been expected by a rational speaker to be optimally relevant to the hearer.

More concisely, the existing sense of the concept of 'correct' may be given as below:

- (13) An overall interpretation is correct if its optimal relevance is rationally foreseeable.

The senses added to these two concepts by the task-specific criterion (11) may be indicated as follows:

- (14) A task-specific interpretation of some aspect of a linguistic utterance - for instance, of its propositional form - is consistent with the principle of relevance, and hence correct, if that task-specific interpretation - for instance, of its propositional form - is instrumental in producing an overall interpretation (of that utterance) which might have been expected by a rational speaker to be optimally relevant to the hearer.

And, more concisely, the added sense of the concept of 'correct' may be given as below:

- (15) A task-specific interpretation is correct if it is instrumental in producing an overall interpretation whose optimal relevance is rationally foreseeable.

It seems that, in terms of the task-specific criterion (11), the criterial property - referred to by Sperber and Wilson (1986:184) in the words "itself ... consistent with the principle of relevance" - is as it were transmitted to the task-specific interpretation from the overall interpretation. It seems too that the connection along which the criterial property is "transmitted" is the instrumentality of the task-specific interpretation in producing the overall interpretation. To put it another way, the task-specific interpretation "acquires" the criterial property by virtue of being *instrumental in producing the overall interpretation*.

In the rest of 4.3 we will try to spell out how all of this works in the case of a hearer faced with the task of identifying the unique propositional form intended by a speaker who has uttered the sentence *Peter's bat is too big*. For ease of reference, we repeat the speaker's utterance here as (16):

- (16) Peter's bat is too big.

4.3.2 Disambiguation

As we noted in 2.3.2 and 2.5 above, the hearer's task involves a subtask of disambiguation. The nature of this task can now be illustrated with reference to utterance (16) and the two semantic representations sketched in (8) and (9). That is: either the hearer must select (8), which incorporates the 'hitting instrument' sense of *bat*, or the hearer must select (9), which incorporates the 'flying mammal' sense of *bat*, as the semantic representation on which to base the rest of his interpretation of this utterance.

In terms of the task-specific criterion (11), then, what is the general procedure for the hearer to use in assigning a semantic representation to an utterance of an ambiguous sentence? In the simplest case, such as that of (16), only two semantic representations are recovered by the linguistic decoding of the utterance. So in the simplest case the general procedure requires the hearer to take at least the four steps sketched below:

(17) *A general disambiguation procedure for the hearer (in its simplest version)*

Step one: given the two semantic representations recovered by the linguistic decoding of the utterance, identify the most readily accessible interpretation which, in order to resolve the ambiguity, assigns to the utterance a specific one of the two semantic representations.

Step two: test whether the interpretation so identified leads to an overall interpretation which is consistent with the principle of relevance - that is, an overall interpretation which might have been expected by a rational speaker to be optimally relevant to the hearer.

Step three: if the task-specific interpretation so tested does lead to such an overall interpretation, assume that this task-specific interpretation is the correct one.

Step four: if the task-specific interpretation so tested does not lead to such an overall interpretation, assign the other semantic interpretation to the utterance.

Of course, if the sentence uttered is ambiguous in more than one way, the linguistic decoding of the utterance recovers three or more semantic representations. To cope with such cases, our simplest version (17) of the general procedure would have to be expanded accordingly. That is, our statement of the general procedure would have to make provision for steps in which the hearer tested alternative task-specific interpretations in decreasing order of accessibility. But as "Step one" of our simplest version (17) makes clear, the notion of 'accessibility of an interpretation' is in any event central to such a general procedure. So, to understand how the general disambiguation procedure even in its simplest conceivable version might work in specific cases, one needs to have some understanding of the intended notion of 'accessibility'. It is to a partial account of this notion that we now turn.

Recall that ostensive stimuli in general, and linguistic utterances in particular, are always interpreted in a context.²¹ In 3.3.2 we presented Sperber and Wilson's notion of 'a cognitive environment': an individual's cognitive environment at a given time is the set of assumptions about the actual world which at that time he is capable of constructing and of accepting as true or probably true. A hearer's context at a given time for interpreting an utterance is always some subset of his cognitive environment at that time or, equivalently, some subset of the assumptions which at that time make up his representation of the world.²² Which assumptions from the hearer's cognitive environment at a given time will be part of his context for interpreting a specific utterance depends heavily on how accessible these assumptions are relative to one another.

The assumptions making up a hearer's cognitive environment at a given time may come from various sources. They may come from long-term or encyclopaedic memory, from short-term memory (including the assumptions which played a role in the interpretation of the immediately preceding utterance), and from concurrent outputs of the other perceptual systems, such as vision.²³ Not all the assumptions making up (the speaker's or) the hearer's representation of the world at a given time are equally accessible as contexts for the interpretation of an utterance. How accessible such assumptions are as contexts depends on various factors; let us note three of these in passing.²⁴ Firstly, the factor of *recency*: the more recently some information has been processed, the more accessible it is. Secondly, the factor of *frequency*: the more frequently some information is used (and so the more stereotyped it has become), the more accessible it is. As for encyclopaedic information, it is organised into blocks filed at conceptual addresses; for instance, if the incoming information being processed by a hearer contains a mapping on to his 'cricket' concept, all the information filed there about cricket becomes accessible to him.²⁵ Sperber and Wilson (1986:138) actually speculate that the encyclopaedic entry of a concept becomes accessible to a speaker or hearer only when that concept appears in an assumption which he or she has already accessed. Moreover, the filed information in each block of encyclopaedic information is organised hierarchically and, as a result, is also ordered in terms of accessibility: the higher up in the hierarchy some information is stored, the more accessible it is. So thirdly - in addition to recency and frequency - *conceptual hierarchy* also affects the accessibility of (the speaker's or) the hearer's assumptions.

The accessibility of assumptions is, in its turn, as we noted in passing in 4.3.1, a significant determinant of the amount of effort required to process an utterance. The more accessible a

²¹ The "situatedness" or context-bound nature of the interpretation of ostensive stimuli is characterised in some detail in ch.3. Cf. in particular 3.3.2 and 3.6.

²² This notion of 'context' was invoked for the first time in our presentation of Case 3.A in 3.3.2.

²³ As we noted in 3.2, S&W (1986:71) see the individual's linguistic system as one of the input systems of his mind along with the five perceptual systems.

²⁴ To date, there is no detailed theory of accessibility. The factors mentioned in the text as determinants of accessibility are among those mentioned by S&W (1986:138ff.) and Carston (1988:61).

²⁵ Sperber and Wilson (1986:86-93) argue for, inter alia, the "speculative" view that concepts are represented in the human mind "as triples of entries, logical, lexical and encyclopaedic, filed at an address" (p.92) and that the encyclopaedic entry attached to a concept "contains information about its extension and/or denotation: the objects, events and/or properties which instantiate it" (p.87).

specific interpretation is relative to alternative interpretations, the less will be the effort required to process it. Conversely, the less accessible an interpretation is relative to alternative interpretations, the greater will be the effort required to process it. Also: the more easily accessible a set of assumptions is, the less will be the effort required to process some new information in a context consisting of those assumptions. And, conversely, the less accessible a set of assumptions is, the greater will be the effort required to process some new information in a context consisting of those assumptions.

Let us now take up again the case of (16), an utterance of the ambiguous sentence *Peter's bat is too big*. Because the sentence uttered is ambiguous, the linguistic decoding of the utterance recovers two different semantic representations for the sentence. We sketched these very roughly, above, as (8) and (9) (and occasionally for convenience will refer to them as (8) and (9), respectively). Now obviously, so far as the hearer's disambiguation task is concerned, two task-specific interpretations are possible: one of these assigns to the utterance the semantic representation (8), and the other assigns to it the semantic representation (9). For convenience we will on occasion refer to them as *the interpretation assigning (8)* and *the interpretation assigning (9)*, respectively. Let us outline now a situation in which one of these task-specific interpretations is clearly more accessible than the other.

Case 4.A

Suppose that the speaker who utters (16) and the hearer to whom (16) is addressed are watching a game of cricket. As they watch the game, they process input received through, especially, their visual and auditory systems. For both speaker and hearer the (visual and auditory) information which they process will contain as a subpart a mapping on to their 'cricket' concept. Consequently, all the information filed at their conceptual addresses for cricket becomes accessible to them, including the 'hitting instrument' sense of *bat*. If Sperber and Wilson are correct to suggest that the encyclopaedic entry of a concept becomes accessible only when that concept figures in an assumption which has already been accessed, the 'flying mammal' sense of *bat* will not be accessible in this situation: there is no reason to assume that in this situation either the speaker or the hearer has processed an assumption containing the 'flying mammal' sense of *bat*. At the very least, therefore, the 'flying mammal' sense of *bat* is much less accessible in this situation than is the 'hitting instrument' sense. To relate this to the two semantic representations of (16) sketched above and to the two possible task-specific interpretations: in the situation just outlined, the interpretation assigning (8) is much more accessible than the interpretation assigning (9).

In terms of the consistency criterion, therefore, the hearer will take as correct the interpretation assigning (8). Accordingly, in Case 4.A, he will assign (8) as the semantic representation on which to base his further interpretive processing of the utterance. In terms of the procedure outlined in (17), however, the assignment of (8) is at best provisional. By the task-specific criterion, the hearer is entitled to accept the assignment of (8) as the correct interpretation only if it leads to an overall interpretation which meets the consistency criterion. In other words: in order to qualify as correct, the task-specific interpretation assigning (8) must be instrumental in producing an overall interpretation which might have been expected by a rational speaker to be optimally relevant to the hearer. Such an overall

interpretation will be produced if, on the one hand, there is an easily accessible context in which the task-specific interpretation assigning (8) has adequate contextual effects for the hearer and if, on the other hand, there are no obviously cheaper means to achieve these effects.

In the cases presented below, we will continue our discussion of possible interpretations of the utterance (16), *Peter's bat is too big*. Our strategy will be to vary deliberately our illustrative situation of speaker and hearer as spectators at a cricket match, and to explore the consequences which each variation has for the interpretive possibilities. Our first variation, presented in Case 4.B, is meant to illustrate one way in which the task-specific interpretation assigning to utterance (16) the semantic representation (8), with *bat* in the sense of 'hitting instrument', might be instrumental in producing for this utterance an overall interpretation which meets the consistency criterion.²⁶

Case 4.B

Suppose that an individual by the name of *Peter*, known to both speaker and hearer, is batting at the moment of the speaker's uttering (16), and that this individual is faring less well than could be expected. Alternatively, one may consider, for instance, a situation in which this individual has just been dismissed after a particularly unsuccessful batting performance. But let us return to our *Peter*, who is still "not out", even though he is performing poorly at the time when (16) is spoken in Case 4.B. In Case 4.B, then, (at least) two assumptions are readily accessible to the hearer as possible contexts for the interpretation assigning the semantic interpretation (8) to the utterance (16). The first of these is the assumption that *Peter* has been batting poorly: an assumption based in part on input from the visual system and in part on items drawn from the hearer's encyclopaedic knowledge of cricket. Let us represent this first contextual assumption as follows:

- (18) *A first readily accessible assumption*
Peter is batting poorly.

The second easily accessible assumption is wholly drawn from the hearer's encyclopaedic knowledge of cricket. Here it is:

- (19) *A second readily accessible assumption*
If a batsman is using a bat which is too big for him, this may cause him to bat poorly.

²⁶ For the purposes of Case 4.B, it will be necessary to assume that certain of the other subtasks that go into interpreting the utterance are performed in ways ending in specific outcomes. In effect, then, when we refer to (8) below we mean a completed (8). For instance, we will assume that the referent of the referential expression *Peter* is the individual batting at the moment of utterance, and that the meaning of the vague phrase *too big* is enriched to 'too big for successful batting'. In all instances the task outcomes adopted here seem intuitively plausible. In later parts of this paragraph it will be shown that each of these task outcomes does in fact meet the task-specific criterion.

These two contextual assumptions can then combine with (a completed) semantic representation (8) as the premisses of a deductive argument which yields the following contextual implication for utterance (16):²⁷

(20) *A contextual implication*

Peter is batting poorly because he is using a (cricket) bat which is too big for him.

So, Case 4.B illustrates that, if the interpretation assigning (8) is processed in the situation as sketched, the hearer is able to derive a probable cause for Peter's poor performance as a cricket batsman. Intuitively, the contextual implication (20), by pinpointing a cause for Peter's poor batting performance, makes it worth the hearer's while to have processed the interpretation assigning (8) in the easily accessible context outlined above. That is, when processed in this cricket-related context, the interpretation assigning the semantic representation (8) to the utterance (16) provides the hearer with *adequate contextual effects* in return for *the minimum of justifiable processing effort*. (This intuition about the adequacy of the contextual effects will be strengthened if, for instance, the speaker and/or the hearer speculated about/commented on Peter's performance at some earlier point during his innings.) As long as a rational speaker could have foreseen this situation, the interpretation assigning (8) meets the task-specific criterion. In terms of the task-specific criterion, the interpretation assigning (8) to the utterance (16) is correct: it is the first accessible task-specific interpretation to lead to an overall interpretation which meets the consistency criterion.

But, in the situation sketched above, is the interpretation assigning (8) the *only* interpretation to meet the task-specific criterion? Recall that in 3.6 we presented Sperber and Wilson's claim that every utterance has *at most one* interpretation which is consistent with the principle of relevance. They claim, in other words, that every utterance has at most a single interpretation which meets the consistency criterion (10). And, as we saw in 4.3.1, they also claim the following: that, operating in tandem with and depending on the consistency criterion, there is the task-specific criterion (11). This additional criterion, in effect, says to the hearer: take as correct that most readily accessible task-specific interpretation which leads to an overall interpretation which meets the consistency criterion (10).

Now, the argument that the interpretation assigning (8) to the utterance (16) is the *only* one to meet the task-specific criterion in the situation as sketched goes as follows:²⁸ imagine that the speaker, in uttering (16), had intended to communicate the semantic representation (9), but that she had foreseen that the interpretation assigning (8) would not only be more accessible than the one assigning (9) but would also meet the task-specific criterion. By picking a more suitable sentence for use in her utterance - perhaps by saying for example *Peter's fruit bat is too big* or *Peter's bat is too big to be vivisected* - she could have spared her hearer a great deal of processing effort. Specifically, she could have spared her hearer the effort of first accessing and processing the interpretation assigning (8), then accessing

²⁷ For a more detailed account of the derivation of implicatures cf. 4.4 below.

²⁸ We take this argument over from W&S (1987a:14-15), where they discuss the interpretation of the ambiguous utterance *The football team gathered round their coach*. The ambiguity of this utterance arises from the fact that an utterance of *coach* can be taken to mean either 'bus' or 'games teacher'. W&S (1986:76) give a similar argument in regard to the ambiguous *I'll bring a bat*.

and processing the interpretation assigning (9) and, finally, engaging in some further form of inference process to decide between these two "competing" interpretations. In other words, by a more judicious choice of utterance she could have achieved the intended range of contextual effects at a much reduced processing cost. To put it differently: in relation to (16), namely an utterance of *Peter's bat is too big*, an utterance of some better-suited sentence would have been an obviously cheaper means for her to achieve the intended contextual effects. In the situation sketched above, the task-specific interpretation assigning (9) is instrumental in producing an overall interpretation which, it is true, is not in principle incapable of achieving an adequate range of contextual effects; equally undeniably, however, the overall interpretation so produced causes the hearer to arrive at these adequate effects in a manner that is roundabout and that therefore puts him to some unjustifiable effort. In the terminology of 3.5: the interpretation assigning (9) leads to an overall interpretation which on the one hand does fulfil the adequate effect presumption of the presumption of optimal relevance, but which on the other hand fails to fulfil the least effort presumption.²⁹ Only the interpretation assigning (8) leads to an overall interpretation which fulfils both these presumptions. Therefore, only (8) is consistent with the principle of relevance in the sense that it meets the task-specific criterion. So: the first accessible task-specific interpretation tested and found to lead to an overall interpretation consistent with the principle of relevance is the only such task-specific interpretation. All other task-specific interpretations will lead to overall interpretations which *fail to fulfil the least effort presumption*, even if they do fulfil the adequate effect presumption.

In the argument outlined above we concentrated on the fact that, in the situation sketched, the assignment of (9) requires more processing effort than does the assignment of (8), and so fails to fulfil the least effort presumption. Let us note that this actually has serious consequences for a speaker who uses (16) with the intention to convey (9), rather than the more accessible (8): there is a danger that the hearer may be satisfied with the contextual effects yielded by (8), and so may not even consider (9). This means, then, that the speaker who uses (16) to convey (9) in the situation sketched above runs a great risk of being misunderstood. If the intended task-specific interpretation is not the correct one in terms of the task-specific criterion, there is then a great likelihood that communication may fail.

So far, we have been considering situations where, of the two interpretations assigning semantic representations to an ambiguous utterance, one is much more accessible than the other. But, of course, ambiguous utterances such as (16) are not produced only in situations in which one of two such task-specific interpretations is more accessible than the other. Let us consider now a second "route" which the interpretation of an ambiguous utterance can follow: a route on which the two interpretations assigning semantic representations are equally accessible, but on which an easily accessible context yielding adequate contextual effects is available for only one of these two interpretations.

²⁹ For the content of the presumption of relevance cf. (15) in 3.5 above.

Case 4.C

In what kind of situation will the two task-specific interpretations assigning (8) and (9), respectively, to the utterance (16) be equally accessible to the hearer? Let us take it that, as sketched before, the speaker and the hearer are watching a cricket game in which a person named *Peter* is batting. But let us imagine that, as they watch the game, speaker and hearer are discussing a biology experiment which involves mammals. More specifically, they are discussing the suitability of a bat (a flying mammal, that is) selected by Peter for the purpose of this experiment. After a lull in the conversation, during which the cricket game has produced nothing especially noteworthy, the speaker utters (16). In this kind of situation, the two interpretations assigning (8) and (9) will probably be equally accessible to the hearer. If so, they will then be tested simultaneously on the task-specific criterion.

Suppose further that the hearer has easy access to a context in which the interpretation assigning (8) has an adequate range of contextual effects, but that a comparable context for the interpretation assigning (9) is much less accessible or not accessible at all. Again, here, we only need to use our imagination. Suppose that Peter is batting poorly as before, so that for the interpretation assigning (8) a context as outlined above (consisting of the assumptions (18) and (19)) is easily accessible. Suppose moreover that the criteria for selecting suitable animals for the biology experiment expressly exclude size as a factor (although they do include age and sex, for instance). Then the hearer does not have easy access to a context in which the interpretation assigning (9) can be tested for contextual effects.

In Case 4.C, the interpretation assigning (8) will again be the first accessible interpretation which meets the task-specific criterion. In line with the procedure outlined above, the hearer will first process the interpretation assigning (8), since only this interpretation has an easily accessible context. This interpretation will have adequate effects in this context, as outlined above. But is it the *only* interpretation to meet the task-specific criterion? Could not the interpretation assigning (9) also meet the task-specific criterion? No, it could not. Not, that is, if a rational speaker could have foreseen that the hearer (i) would have easy access to a context in which the interpretation assigning (8) had adequate contextual effects, and so (ii) would process this interpretation first and judge it to meet the task-specific criterion. If a rational speaker had wanted to convey the semantic representation (9), she would again have spared the hearer some processing effort by reformulating her utterance appropriately - that is, by picking for her utterance a sentence whose semantic representation does not lend itself to the unwanted interpretation. That is, if a rational speaker had intended to communicate the semantic representation (9), she could have foreseen that the utterance (16) would put the hearer to some unjustifiable processing effort - unjustifiable, because, as we have just explained, an obviously easier way was open to the speaker for communicating the semantic representation she intended. In the hypothetical situation outlined here, the interpretation assigning (9) would lead to an overall interpretation which failed to fulfil the least effort presumption, and which consequently failed to meet the consistency criterion. (Whether it meets the adequate effect presumption is beside the point.) Again, then, the first accessible interpretation tested on the task-specific criterion and found to meet it, and found in this way to be consistent with the principle of relevance, is the only interpretation consistent with the principle of relevance.

There is also a third possibility for the processing of an ambiguous utterance. It is possible for the two interpretations assigning the alternative semantic representations to an ambiguous utterance to be equally accessible and to yield, moreover, comparable contextual effects for comparable processing costs. Let us see if this possibility too can be made concrete with reference to the task-specific interpretation assigning the semantic representations (8) and (9), respectively, to the utterance (16).

Case 4.D

Suppose that (16) is uttered in a situation which, except for one difference, is identical to the situation sketched in Case 4.B. Here is the difference: the criteria for selecting a suitable mammal for the biology experiment expressly *include* the size factor. The interpretation assigning (8) will then yield adequate contextual effects in the context of the assumptions (18) and (19). It is not too hard to imagine some readily accessible assumptions with which the interpretation assigning (9) could combine to yield adequate contextual effects. For instance, let us take it that this interpretation combines with the assumption that there are only two possible subjects for the experiment, namely Peter's bat and a certain mouse. Then the outcome is the implication that the mouse must be selected for use in the experiment.

In Case 4.D the hearer will have no way of choosing between the two interpretations assigning (8) and (9), respectively, and so the ambiguity will remain unresolved. In this case, Sperber and Wilson argue, *neither* of the two interpretations meets the consistency criterion: neither could be preferred, if at all, until after an effort of comparison which a rational speaker could have spared the hearer, if she could have foreseen that the two interpretations would be equally accessible and would yield comparable contextual effects at comparable processing costs. That is, for both these interpretations there is an obviously cheaper way in which the speaker could have achieved her intended contextual effects. As Wilson and Sperber (1987a:15) note, such cases illustrate that the principle of relevance provides an account not only of *successful* disambiguation, but also of *failures* of disambiguation.³⁰

4.3.3 Reference assignment

In 4.3.2 we illustrated how the principle of relevance is implemented by means of the task-specific criterion, operating in tandem with the consistency criterion. In particular, we illustrated how the hearer of an ambiguous utterance, given the semantic representations recovered for the utterance by linguistic decoding, is guided by the principle of relevance to select one of these as the basis for his overall interpretation of the utterance. The task-specific criterion also guides the hearer in expanding the selected semantic representation to a propositional form (that is, a semantically complete logical form). In 4.3.3 we will focus

³⁰ In this account of how the task-specific criterion guides the hearer in disambiguating an ambiguous utterance, we rely heavily on the sources referred to in n.28. The disambiguation task is also discussed by S&W (1986:185-187); they suggest there that the relation between the decoding process carried out by the hearer's linguistic module and the inferential process carried out by the hearer's central processing mechanisms may well be more complicated than is apparent from the discussion above. Nevertheless, our account is faithful to their basic claim: that disambiguation

on how the task-specific criterion enables the hearer to construct and evaluate inferences, or, equivalently, hypotheses, about the referents of all the indeterminate referring expressions which occur in the selected semantic representation.³¹ In 4.3.4 we will focus on how the task-specific criterion guides the hearer in performing the task of enriching any semantically incomplete terms which occur there.

Given the task-specific criterion, the hearer faced with the task of assigning referents to indeterminate referring expressions must choose the solution which requires the least effort, that is, the most accessible solution. He should abandon this solution only if it fails to lead to an overall interpretation which a rational speaker could have expected to be optimally relevant to the hearer. Sperber and Wilson (1986:187) give the following statement of how this will work for a hearer faced with the task of interpreting the referentially indeterminate expression *it* which occurs in the utterance *It will get cold*:

"Given the principle of relevance, he should first consider the immediate context, see if any of the concepts of a non-human entity represented in this context, when substituted for 'It', yields a propositional form consistent with the principle of relevance; if not, he should extend the context and repeat the procedure."

Let us see how this applies to the interpretation of the referentially indeterminate expression *Peter* which occurs in (16). Let us consider again the situation sketched as Case 4.B. In this situation, as we have seen, the hearer will select the semantic representation (8) as the basis of his further interpretation of the utterance. (8) is not a propositional form, in part because it contains the referentially indeterminate expression *Peter*. In terms of the task-specific criterion, the hearer faced with the task of expanding (8) into a propositional form should first consider the immediate context, in order to determine whether any of the concepts of a human entity available in this context is identifiable by the label *Peter*. Provided that the hearer knows that the cricket player who is batting so poorly is called *Peter*, the most accessible hypothesis – that is, the one requiring the least effort to construct – regarding the possible referent of the expression *Peter* will be that this referent is the player currently batting so poorly. In order to evaluate this hypothesis, the hearer must determine whether the inference that this batsman is the intended referent leads to an overall interpretation which meets the consistency criterion. That is, he is to check whether it leads to an overall interpretation which a rational speaker might have expected to be optimally relevant to the hearer. In Case 4.B, as outlined in 4.3.2, this inference will indeed lead to such an overall interpretation. The correct task-specific interpretation assigning a reference to the expression *Peter* in the situation outlined in Case 4.B is, then, that *Peter* refers to the cricket player who is batting so poorly at the time of utterance.

In this situation, is there any other interpretation for the expression *Peter* that will meet the task-specific criterion? The answer is that, as with the resolution of ambiguity, there is at most one interpretation for a referentially indeterminate expression that meets the task-

hypotheses are recovered by decoding and are evaluated inferentially, under the guidance of the task-specific criterion.

³¹ For a discussion of how the consistency criterion guides the hearer in his attempt to determine the intended referents of referring expressions cf. S&W 1986:187-188.

specific criterion. We can show this by constructing an argument parallel to the one constructed above to show that at most one of the semantic representations for an ambiguous utterance is consistent with the principle of relevance. Consider once more the situation outlined as Case 4.B above. Imagine that the speaker, in uttering (16), had wanted to refer not to the person bearing the name *Peter* and currently batting, but to some other batsman called *Peter*, not involved in the game she and the hearer are watching. Imagine also that the speaker had foreseen that the hypothesis that *Peter* refers to a person bearing the name *Peter* and currently batting would be both more accessible and consistent with the principle of relevance. Then she could have eliminated this unwanted interpretation by reformulating her utterance – for instance, by adding the surname of the *Peter* to whom she had intended to refer. In this way she could have spared her hearer some unnecessary processing effort. Although the alternative hypothesis that *Peter* refers to some other person bearing the name *Peter* but not playing in the current game might lead to an overall interpretation which fulfilled the adequate effect presumption, such an overall interpretation would fail to fulfil the least effort presumption: to a speaker intending to refer to such an individual an obviously cheaper way was open for achieving her intended contextual effects. Therefore, of the above-mentioned interpretations assigning a reference to *Peter*, only the one we outlined first – namely, that *Peter* refers to the person bearing the name *Peter* and currently batting – meets the task-specific criterion.

The discussion above contains an illustration of how the principle of relevance, implemented by means of the task-specific criterion, enables a hearer to carry out the task of identifying the intended referent of a referentially indeterminate expression which occurs in a selected semantic representation. To close our discussion of the task of assigning referents to indeterminate referring expressions, let us have a brief look at a case in which reference assignment is likely to fail.

Case 4.E

Consider once more the situation sketched as 4.B above, but now modified as follows: two persons bearing the name *Peter* are batting, neither of them very successfully. This situation gives rise, of course, to two equally accessible interpretations assigning references to the indeterminate expression *Peter*. Moreover, for each of these task-specific interpretations there is an easily accessible context in which it has adequate contextual effects: in the case of either batsman, the assumption that his bat is too big yields an explanation for his poor batting performance.

In brief: in Case 4.E two interpretations assigning references to *Peter* are equally accessible, and lead to comparable contextual effects at comparable processing costs. The hearer has no way of choosing between these two interpretations, since they both fail to meet the task-specific criterion. In both cases, the speaker could have conveyed her intended reference in an obviously cheaper way by adding to *Peter* something distinctive, such as a surname or a nickname, to enable the hearer to eliminate the unwanted interpretation.

This brief discussion of Case 4.E illustrates that in the case of reference assignment, too, the task-specific criterion in conjunction with the consistency criterion is able to explain not only successes, but also failures.

4.3.4 Enrichment of vague terms

As we noted in 4.2, the semantic representations of (16) contain two constituent-meanings which are semantically incomplete in that they are vague or "underdefined" by the generative grammar of the language: the genitive 's in *Peter's bat* has the vague meaning 'somehow associated with' (= (4) above), and *too* in *is too big* has the vague meaning 'excessive relative to some unidentified reference-point' (= part of (7) above). As a result, the semantic representations (8) and (9) of (16) are incomplete in two respects: (i) they fail to specify the exact nature of the relationship between the respective meanings of *Peter* and *bat*, and (ii) they fail to specify the reference-point or standard relative to which the "bat" is "too big". The third subtask which a hearer of (16) must carry out in order to assign this utterance a unique propositional form is, then, to enrich or, equivalently, complete at these points whichever of the two semantic representations he selects.

Sperber and Wilson (1986:189) argue that the hearer's ability to supply these enrichments is also explained by the task-specific criterion (11). If their argument is sound, this criterion enables the hearer to identify the correct enrichment for a vague constituent-meaning from its being the first accessible enrichment that leads to an overall interpretation which meets the consistency criterion (10) - that is, an overall interpretation which might have been expected by a rational speaker to be optimally relevant to the hearer. And so, if this argument of Sperber and Wilson's is sound, the hearer, in searching for the correct interpretation enriching a vague expression, is guided by the same additional criterion which enables him to carry out the disambiguation task and the reference-assignment task. Here, as there, the hearer must select the first accessible task-specific interpretation. Here, as there, if the hearer finds that this task-specific interpretation leads to an overall interpretation which meets the consistency criterion, he must assume that this task-specific interpretation is indeed the one the speaker intends.

For the purpose of our illustration, let us take up again the situation sketched in Case 4.B above. And let us take it up at the stage where the hearer has selected the semantic representation (8) - incorporating the 'hitting instrument' sense of *bat* - as the basis for his further interpretation, and has inferred that *Peter* refers to the batsman who is batting so poorly at the time of utterance. In such a situation, the first accessible interpretation enriching the vague expression 's in *Peter's bat* is likely to enrich it from 'somehow associated with' to 'being used by'; accordingly, *Peter's bat* is then interpreted to mean (very roughly) 'the hitting instrument being used by the person who bears the name *Peter* and who is batting at the time of utterance'. The intrinsic meaning of the predicate phrase *is too big* is also easy to enrich in the situation outlined in Case 4.B, namely to (very roughly) 'too big for Peter to use for successful batting'. Of course, each of these first accessible task-specific interpretations enriching vague terms needs to be tested; as with any task-specific interpretation, the hearer has to check whether it leads to an overall interpretation of (16) which meets the consistency criterion. In our discussion of Case 4.B we suggested such an overall interpretation for (16) - namely, an interpretation which leads the hearer to infer a probable cause for Peter's poor batting performance. Such an overall interpretation fulfils not only the adequate effect presumption but also the least effort presumption. Consequently, the hearer will infer that his first task-specific interpretations enriching the two vague expressions - *Peter's bat* and *too big* - are the intended interpretations.

Again one can construct an argument to the effect that these are the *only* possible task-specific interpretations. The argument again relies on the assumption that no other task-specific interpretation would lead to an overall interpretation which met the least effort presumption of the presumption of optimal relevance. Any other task-specific interpretation would cause the hearer to use more than the minimum of justifiable processing effort. If the speaker had intended any other interpretation, and she could rationally have foreseen that the task-specific interpretations outlined above would be the first to be tested and found consistent with the task-specific criterion, she could have spared her hearer such unjustified processing effort by using a more appropriate utterance - that is, by uttering a sentence whose semantic representation would not give rise to the unwanted interpretation. Again, then, for each of the semantically vague expressions there is at most a single interpretation which meets the task-specific criterion.

Let us make these last points concrete by focusing on the interpretation of the expression *too big*. Again, consider the situation outlined as Case 4.B above. Suppose the speaker intends to convey to the hearer that Peter's bat is too big to fit into Peter's bag. Let us further assume that in the time before the speaker spoke (16) there was no mention either of bags or of the possibility of a bat's fitting into a bag. A rational speaker will foresee the following: firstly, that in the situation as outlined the interpretation enriching the meaning of *too big* to 'too big for Peter to use for successful batting' will be more accessible than the interpretation enriching it to 'too big to fit into Peter's bag' and, secondly, that this more accessible task-specific interpretation will lead to an overall interpretation which meets the consistency criterion. If the speaker wants to convey the interpretation 'too big to fit into Peter's bag', she can spare her hearer some unjustified processing effort, and cut down the risk of being misunderstood, by picking for her utterance some more suitable sentence, for instance *Peter's bat is too big to fit into his bag*.³²

4.3.5 Summary

Recall that in order to bridge the gap between the set of incomplete logical forms contained in the semantic representations of an utterance and a unique, propositional (logical) form for this utterance, the hearer must perform three subtasks: (i) disambiguate the utterance, that is, select one of the semantic representations as the basis for the interpretation of the utterance; (ii) assign referents to all the referentially indeterminate terms in the selected semantic representation; (iii) complete the selected representation by sufficiently enriching all vague, semantically incomplete, terms. Which of the many possible propositional forms is the correct one - that is, the one intended by the speaker - is determined by the task-specific criterion (11): the correct propositional form is that form which leads to an overall interpretation which meets the consistency criterion (10). To be more explicit: an overall interpretation meets the consistency criterion if it might have been expected by a rational

³² Commenting on some differences between their own views and those of Grice and of pragmatists influenced by him, S&W (1986:182-183) assert that "[m]ost Gricean pragmatists assume without question that any pragmatically determined aspect of utterance interpretation apart from disambiguation and reference assignment is necessarily an implicature." In terms of the assumption attributed here to "most Gricean pragmatists", the kind of enrichment of vague expressions which is illustrated in 4.3.4 would presumably be a pragmatically determined aspect of the implicit content of utterances and not, as for S&W, of the explicit content of utterances.

speaker to fulfil the presumption of optimal relevance. To be more explicit still: an overall interpretation fulfils the presumption of optimal relevance if it meets both the adequate effect presumption and the least effort presumption of the presumption of relevance ($\approx$ (15) in 3.5 above).

In the process of constructing the unique propositional form intended by the speaker, the hearer is guided by the task-specific criterion (11), which operates in conjunction with the consistency criterion. This guidance can be said to comprise (at least) three items. First item: for every task, the hearer should identify the first accessible interpretation (that is, the one requiring the least effort to construct). Second item: the hearer should check whether this first accessible interpretation leads to an overall interpretation which meets the consistency criterion. Third item: only if this first accessible task-specific interpretation fails to lead to an overall interpretation which meets the consistency criterion, must the hearer consider a more costly alternative.

In our illustrations we have tried to make these general points concrete by showing what propositional form the hearer is likely to construct for (16) in the situation sketched in Case 4.B and in certain variants of this situation. To recap briefly the various points made in connection with the interpretation of (16) in Case 4.B: (i) the hearer selects the semantic representation (8) as the basis for his further interpretation of (16); (ii) the hearer infers that the intended referent of the referentially indeterminate expression *Peter* is the person who bears the name *Peter* and who is batting at the time of utterance; (iii) the hearer infers that the vague expression *Peter's bat* is intended to convey the idea 'the hitting instrument used at the time of utterance by Peter, who is the batsman currently batting', and that the vague expression *too big* must be interpreted as 'too big for Peter to use for successful batting'. By successfully completing these various tasks, the hearer recovers a propositional form whose content may be informally represented as below:

(21) *The propositional form of utterance (16): a rough sketch*

'At the time of utterance the following holds: the hitting instrument being used by the person who bears the name *Peter* and who is batting has the property of being too big for him to bat with successfully.'

As is clear, we hope, from our illustration of the recovery of this propositional form, the recovery of the propositional forms of utterances is an inferential process guided by the criterion of consistency with the principle of relevance. As with all non-demonstrative inferential processes, so here too the fallibility of non-demonstrative inference asserts itself:³³

(22) *The fallibility of non-demonstrative inference*

There can be no guarantee of the correctness of non-demonstrative inference.

³³ This fallibility is due to a property which we pointed out in n.9 to ch.3: non-demonstrative inference or, equivalently, every non-demonstrative argument "involves the claim, not that its premisses give conclusive evidence for the truth of its conclusion, but only that they provide *some* evidence for it" (cf. Copi 1973:21).

That is, the consistency criterion guides the hearer of an utterance to the *best* hypothesis about its propositional form, but it does not guarantee that the best hypothesis is the *correct* hypothesis (Sperber and Wilson 1986:65-71; Wilson and Sperber 1986:70-71; 1987a:8).

In the situation sketched in Case 4.B, for instance, the speaker's thoughts, without the hearer's being aware of it, may have digressed to the biology experiment referred to under Case 4.D. The speaker may moreover become quite engrossed in her thoughts about this experiment. As a result, she may then utter (16) to convey to the hearer certain information relating to this experiment. This may be the information that the *flying mammal* offered by a mutually known individual, Peter (possibly none other than the Peter batting at that moment), for use in this experiment is too big for that purpose. And the speaker, in addressing (16) to the hearer, may do so with this informative intention without considering that he has no way of knowing that the situation which she is referring to is that of the experiment. Again, the hearer will consider first the "cricket" propositional form. If this propositional form turns out to yield adequate contextual effects, as outlined above, then the hearer will probably misunderstand the utterance. Clearly the case in which he *does* (misunderstand the utterance) is a case in which the best hypothesis about the content of the propositional form of the utterance is not the correct hypothesis. Communication, in such a case, will have failed.

As we emphasised at the start of 4.3.1, recovering the propositional form of an utterance is the same thing as recovering its explicit content. Accordingly, the informal representation (21) above can be said to represent, either the propositional form, or equivalently the explicit content, which the utterance (16) conveys when spoken in the situation we outlined in Case 4.B above. As we showed in 2.4, however, the overall interpretation of utterances involves the recovery of their implicit content as well. In the following sections, we will describe and illustrate the role which the criterion of consistency with the principle of relevance is claimed to play, by Sperber and Wilson, in the hearer's recovery of the implicit content of utterances.

4.4 How the consistency criterion explains the recovery of the implicatures of ordinary assertions

Ordinary assertions, as conceived of by Sperber and Wilson (1986: 181, 183, 193), have at least two properties which we may note here to start with. Firstly, an ordinary assertion is an utterance which communicates its propositional form explicitly. In this respect, the utterance *Peter's bat is too big* - at least when produced in any of the contexts outlined in 4.3 - is an example of an ordinary assertion: the propositional form of this utterance is part of the content explicitly conveyed by this utterance.³⁴ Secondly, an ordinary assertion is an utterance most of whose contextual effects, and therefore most of whose relevance, depend on its propositional form.³⁵ In the case of *Peter's bat is too big*, again, the relevance of the

³⁴ Ordinary assertions differ from, for instance, ironical utterances: in the case of an ironical utterance the speaker does not intend to communicate the utterance's propositional form explicitly. In this connection cf. the discussion of the ironical utterance (10) in 2.4.5. We return to ironical utterances in 4.7 below.

³⁵ Our characterisation of an ordinary assertion is a slightly simplified version of S&W's (1986:193). In their characterisation, they make use of the notion of 'explicature'. An explicature,

utterance depends largely on its propositional form, as we will show below. In this respect too, then, this utterance is an example of an ordinary assertion.

But of course ordinary assertions do not only convey assumptions explicitly; they can also convey assumptions implicitly. The standard term for such implicitly communicated assumptions is *implicatures*. In this section, we focus on Sperber and Wilson's (1986:194) claim that the recovery of the implicatures of ordinary assertions, like the recovery of their explicit content, is governed by the criterion of consistency with the principle of relevance. The essence of their position, stated informally, is as follows: what makes it possible for a hearer to identify the implicatures of an ordinary assertion is that a rational speaker must have expected the hearer to recover/derive them, or at least some of them, given that she intended her utterance to be optimally relevant to the hearer.³⁶ Let us try now to make this idea more precise. We start off by yet again looking at the interpretation of the utterance (16) - *Peter's bat is too big*.

Consider once more the interpretation of (16) in the situation outlined as Case 4.B above. We illustrated in 4.3 how the task-specific criterion (11), in conjunction with the consistency criterion (10), guides the hearer to construct (21) as the complete, and unique, propositional form for (16) in this situation. In each instance a task-specific interpretation, provisionally picked because of its superior accessibility, has to be tested with regard to the overall interpretation of (16): the provisionally picked task-specific interpretation must be found to be instrumental in producing an overall interpretation which meets the consistency criterion (10). But what is this "overall interpretation"? In discussing Case 4.B, we have suggested that the overall interpretation of (16) includes the assumption presented in (21) - namely, that Peter is batting poorly because his bat is too big. But (20) does not form part of the explicit content of (16), informally presented in (20). Rather, (20) is part of the *implicit* content of (16). This means, then, that, the set of assumptions conveyed explicitly is only a subset of the set of assumptions {*I*} which makes up the intended "overall" interpretation of (16). Two questions arise at this point, the second of these questions presupposing a positive answer to the first one:

- * Firstly, then, there is a *whether* question: how can the hearer of the utterance know that the set of assumptions {*I*} which the speaker intends to communicate contains assumptions in addition to those communicated explicitly by its propositional form (20)? (This, of course, is merely one instance of a more general question: how can the hearer of any utterance tell whether the speaker means it to convey (one or more) implicit assumptions?)
- * Secondly, there is a *which* question: how can the hearer know that what the speaker intends him to recover from the utterance, in addition to its explicit propositional content, is specifically (20) rather than any other assumption? (Again, this is an instance of a more general question: how can the hearer of an

according to S&W (1986: 182), is an explicitly communicated assumption. Explicatures therefore are in contrast to implicitly communicated assumptions, which S&W refer to by the familiar term *implicatures*. According to them (1986:193), then, in an ordinary assertion "the propositional form is itself an explicature, and indeed the explicature on which most of the contextual effects of the utterance, and therefore most of its relevance, depend".

³⁶ For a detailed discussion of how the criterion of consistency with the principle of relevance guides the recovery of the implicatures of ordinary assertions cf. S&W 1986:193-202.

utterance tell which implicit assumption or assumptions the speaker means it to convey?)

In order to understand Sperber and Wilson's answers to these questions in regard to utterance (16), we must consider again how the task-specific criterion (11), in tandem with the consistency criterion (10), enables the hearer of an utterance to recover its intended propositional form.

The task of determining the propositional form of an utterance is of course made up of a variety of subtasks. For the moment, however, let us abstract away from that fact. Then we can say simply that (21) is the first accessible propositional form to be tested by the hearer for consistency with the principle of relevance. The selection of this propositional form as the first to be so evaluated is determined by the task-specific criterion (11). The hearer must then determine whether the propositional form so selected is instrumental in producing an overall interpretation which meets the consistency criterion (10) - an overall interpretation, that is, which is consistent with the principle of relevance. An interpretation is consistent with the principle of relevance if it might have been expected by a rational speaker to be optimally relevant to the hearer. The presumption of optimal relevance communicated by a linguistic utterance has two parts, as was first noted in 3.5-3.6 above and as has been illustrated in 4.3. On the one hand, there is the adequate effect presumption. On the other hand, there is the least effort presumption.

We focus first on the adequate effect presumption. In terms of this presumption, the set of assumptions $\{I\}$ which the speaker intends to make manifest to the hearer is relevant enough to make it worth the hearer's while to process the utterance. This then provides us with an answer to the *whether* question raised above: how can the hearer know whether the speaker intends her utterance to convey (one or more) implicit assumptions? The hearer will begin to search for additional, implicit, assumptions if the explicit content of the utterance does not provide him with adequate effects to offset the effort required for the utterance's processing. In the case of our example, then, the hearer has to begin to search for implicit assumptions if the explicit content of (16) does not in itself provide him with adequate effects to offset the effort he put into (16)'s processing. Suppose, for instance, that the explicit content of (16) does not serve to strengthen or contradict some assumption about the size of Peter's bat already contained in the hearer's cognitive environment. Then the explicit content is not relevant enough in its own right, and the hearer has to begin to search for contextual implications.³⁷

We come now to Sperber and Wilson's answer to the second question raised above: how can the speaker know *which* implicit assumptions the speaker intends to communicate? The hearer, they claim, must recover those assumptions which might have been expected by a rational speaker to make the utterance optimally relevant to him. In his search for these assumptions, the hearer is again guided by the task-specific criterion (11). That is, he should first identify the most readily accessible task-specific interpretation - in this instance, the most readily accessible set of implicit assumptions - and then test whether it leads to an

³⁷ Here we are drawing on S&W's (1986:121) view that having some contextual effect in a given context is a necessary condition for an assumption's relevance in that context. The three types of contextual effect which assumptions can have were illustrated in 3.3.2 above.

overall interpretation which meets the consistency criterion. If this first accessible task-specific interpretation yields adequate effects, the hearer must assume that it is indeed the correct, intended, interpretation. Only if the first accessible task-specific interpretation does not yield adequate effects, should the hearer continue the processing by identifying the next most readily accessible set of assumptions and testing this for compliance with the consistency criterion (10).

The accessibility of the task-specific interpretation in this subtask depends critically on the relative accessibility of the various possible contexts in which the utterance can be processed. The hearer should select the first accessible context - that is, the context whose retrieval requires the least processing effort.³⁸ He must then process the utterance in this first accessible context in order to determine whether its effects in this context are adequate. If it has such effects, and if a rational speaker might have expected him to derive such effects in the context selected, then that interpretation is consistent with the principle of relevance. The hearer may then assume that that is the correct, intended, interpretation, since at most one interpretation of an utterance can meet the consistency criterion (10).³⁹

But let us return to our example. It seems clear that, in Case 4.B, a rational speaker might have expected her hearer not only to recover the assumptions contained in the propositional form (21), but also to derive some assumptions in addition to those, in order for the utterance (16) to be optimally relevant to the hearer. A hearer in the situation outlined as Case 4.B might, moreover, have been expected by a rational speaker to take the following as the most readily accessible context (note that it consists of two assumptions):

(23) *The most readily accessible context*

- (a) Peter is batting poorly.
- (b) If a cricket player uses too big a bat, this may cause him to bat poorly.

Contextual assumptions such as these, Sperber and Wilson (1986:195) call *implicated premisses*.⁴⁰ Implicated premisses have to be supplied by the hearer (for instance, he may have to retrieve them from memory).

³⁸ S&W (1986:132-151) extensively discuss the selection of contexts for the interpretation of ostensive stimuli, which of course include linguistic utterances. Cf. in particular their (1986:142ff.) discussion of how the selection of contexts is affected by accessibility.

³⁹ In addition to the accessibility of the contextual assumptions, there is of course a second factor which affects the processing effort needed to recover a set of implicit assumptions, namely the effort needed to process an assumption in a given context (S&W 1986:144). In their discussion of specific instances of utterance interpretation, S&W do not consider this second factor in any detail. Understandably therefore their discussion does not make it clear how these two factors jointly affect overall processing effort. We can do no more than point out to the reader that, in addition to the accessibility of the assumptions constituting a particular context, the processing of some new assumption in that context may also affect the effort required to recover implicit assumptions, and hence the optimal relevance of the interpretation recovered. In the text we will, like S&W, focus on how processing effort is affected by the accessibility of contextual assumptions.

⁴⁰ As we pointed out in n.10 to ch.2, we use the phrase *a minor premiss* to refer to an assumption/statement such as (23a) - expressed by means of a syntactically simple declarative sentence - and the phrase *a major premiss* to refer to a conditional assumption/statement - expressed by means of an *if, then* sentence such as (23b).

In Case 4.B, then, these implicated premisses provide the hearer with the most readily accessible context. In this context the hearer might also have been expected by a rational speaker to derive the contextual implication (20), statable as *Peter is batting poorly because he is using too big a bat*. To derive this contextual implication, the hearer makes use of an appropriate deductive argument. The deductive argument in which the various assumptions referred to above combine to yield this contextual implication can be informally presented as follows:

- (24) (a) If a cricket player uses too big a bat, this may cause him to bat poorly (= (23b)).
(b) Peter is batting poorly (= (23a)).
(c) Peter is using too big a bat.
(d) Therefore, the fact that Peter is using too big a bat may be what is causing him to bat so poorly (= (20)).

Given the informal level of our exposition, the wording in (d) of this argument, ... *the fact that Peter is using too big a bat may be what is causing him to bat so poorly*, may with impunity be taken as equivalent to *Peter is batting poorly because he is using a (cricket) bat which is too big for him* (= (20)). So (24d), the conclusion of the argument, represents the contextual implication which the hearer derives. (24c), a minor premiss of the argument, represents the explicit content of (16), the utterance which the hearer is interpreting. By means of the argument, he derives the contextual implication from the explicit content of the utterance in combination with the implicated premisses: in this argument, (24b) represents the implicated - minor - premiss (23a), and (24a) represents the implicated - major - premiss (23b). An implication such as (24d) - derived from the implicated premisses and from the explicit content of the utterance - Sperber and Wilson (1986:195) call *an implicated conclusion*. And they use the term *implicatures* as an umbrella term for referring both to implicated premisses and to implicated conclusions.

In Case 4.B, then, what makes it worth the hearer's while to process the utterance (16) is the recovery of an implicated conclusion specifying a possible cause for Peter's poor batting performance. As has been outlined above, the hearer might have been expected by a rational speaker to derive exactly this interpretation. The hearer, as we have seen, will recover this interpretation, if he is guided by the task-specific criterion (11) and by the consistency criterion (10) on which it depends.

To conclude our discussion of the recovery of the implicit content of (16), let us review in brief the argument that, in the situation assumed above (and first outlined as Case 4.B), there is only one overall interpretation which meets the consistency criterion. This argument will be familiar to the reader; here is its major premiss (that is, the premiss expressing an *if, then* relation):

- (25) If a rational speaker *had considered* whether to use the utterance (16) for the purpose of communicating some set of implicit assumptions other than (24d), but in so considering *could have foreseen*
- * that the first context to be accessed by the hearer in interpreting (16) would be the context indicated above (= (23)), and

- * that the hearer would derive adequate contextual effects from (16) combined with that first accessible context,

then a cheaper means than the use of (16) was available to the speaker for achieving her purpose.

Consequently, the speaker could have spared her hearer some unjustifiable processing effort by, for instance, so reformulating her utterance as to make sure that the context she had in mind would be the first one to occur to the mind of the hearer. In general therefore, if the recovery of the intended one of two possible interpretations of an utterance required the hearer to select some context other than the one most accessible to him, then this most accessible context, even though it yielded adequate contextual effects, *would fail to fulfil the least effort presumption* of the presumption of optimal relevance.

Next, we turn to another case where interpreting an ordinary assertion includes the recovery of implicatures. In 2.4.2, in first noting the subtask of determining the implicatures of ordinary assertions, we alluded to the following question: how is the hearer, John, to take the assertion in (26b), when this is spoken in response to the question (26a)?

- (26) (a) *John*: Do you want some coffee?
(b) *Mary*: Coffee would keep me awake.

We can now look at a relevance-theoretic explanation of how the hearer interprets this assertion. For the purpose of illustration, we will outline two situations in which this exchange may take place.

Case 4.F

Let us assume for the sake of argument that John and Mary are husband and wife, at home in an ordinary domestic situation, and not for instance in a supermarket buying groceries. In this domestic situation, the propositional form of Mary's response is, in essence, the information that if the speaker were to have a drink of coffee, this would cause her to remain awake. Informally and very roughly, let us represent this propositional form as follows:⁴¹

- (27) *The propositional form of utterance (26b)*
Drinking some coffee would keep Mary awake.

The recovery of this propositional form may seem quite straightforward. Note, however, that it involves more than just the recovery of the semantic representation of the underlying sentence, *Coffee would keep me awake*. For instance, this semantic representation contains no information about how the entities called "coffee" and "Mary" are to interrelate in order

⁴¹ As 4.3 above is meant to make clear, the recovery of the propositional form of an utterance like (26b) is far more complex than we show it here. In the text we comment briefly on just one aspect of this recovery, namely the enrichment of the meaning of *coffee*. Also, (27) is only a very rough approximation of the content of this propositional form. It is sufficient, however, for the purpose of the present illustration.

for the coffee to cause the effect described. That is, the information in the semantic representation leaves wide open the question of what it would mean for the speaker to "have coffee". Part of the hearer's task is therefore to enrich the semantic representation by "filling in" the appropriate Mary-to-coffee relation. The hearer is able to fill in this relation from his stereotyped knowledge of the world, specifically his knowledge that coffee is usually taken in the form of a beverage. Without this particular item of knowledge, let us note, the hearer would be quite simply unable to work out the propositional form the speaker intends. Fortunately, however, such stereotyped information is highly accessible. The corresponding interpretation - that the speaker in saying "coffee" is referring to the drinking of a coffee beverage - is likewise highly accessible, and therefore requires little processing effort. As is not unreasonable, let us assume that this is in fact the most readily accessible interpretation enriching *coffee* here. In terms of the task-specific criterion (11), moreover, this first accessible interpretation is to be rejected only if it turns out not to lead to an overall interpretation which meets the consistency criterion (10).

Even with *coffee* enriched in this way, however, the propositional form of the utterance contains no information about the speaker's desire or non-desire to drink coffee. That is, no information about the speaker's desire or non-desire to drink coffee is conveyed explicitly by the utterance in question. Nevertheless, intuitively, the utterance will normally be interpreted as conveying just such information. Such information belongs therefore to the utterance's implicit content; more specifically, it constitutes (part of) the utterance's implicatures. How, then, is John to know *that* Mary intends to communicate some information to him *implicitly*? The answer is made possible by the consistency criterion: the utterance communicates to the hearer the presumption that it is optimally relevant to him; yet, by itself, the utterance's explicit content does not fulfil the adequate effect presumption. To be more specific: for the utterance to be optimally relevant to the hearer in the situation outlined above, he must interpret it as conveying information on whether or not Mary wants any coffee. The hearer, John, therefore has to search for implicitly conveyed assumptions, in addition to the explicitly conveyed assumptions, in order to recover a set of assumptions $\{I\}$ whose optimal relevance to the hearer might have been foreseen by a rational speaker. How then is John to know exactly *which* further implicit *assumption(s)* to add to the set of explicitly conveyed assumptions? Again, it is the task-specific criterion (11) that guides the hearer in performing this subtask of utterance interpretation.

Suppose the contextual assumption that Mary does *not* want to take anything that will keep her awake is more accessible to John than the contextual assumption that Mary *does* want to take something that will keep her awake. In what kind of situation will there be this difference in accessibility between these two contextual assumptions? Let us mention just two such situations. Firstly, there will be this difference in accessibility if it is mutually manifest to speaker and hearer that Mary regularly has difficulty in sleeping. Secondly, there will be this difference in accessibility if Mary has indicated to the hearer at some earlier point in their interaction that she is looking forward to a good night's sleep. So, the more readily accessible of the two contextual assumptions can be presented as below:

- (28) *The most readily accessible context*
Mary does not want to stay awake.

John can then combine this contextual assumption and the propositional form of the utterance to derive the contextual implication that Mary does not want any coffee. He does so by means of the argument which can be informally outlined as follows:

- (29) (a) If Mary does not want to stay awake, and if drinking some coffee would keep Mary awake, she will not want to drink any coffee.
- (b) Mary does not want to stay awake.
- (c) Drinking some coffee would keep Mary awake.
- (d) Therefore, Mary does not want to drink any coffee.

The conclusion (29d) of this argument constitutes an implicature of the utterance (26b). For convenience, we explicitly present this implicature as follows:

- (30) *An implicature recovered for utterance (26b)*
Mary does not want to drink any coffee.

Again, then, the first accessible interpretation tested and found to be consistent with the principle of relevance is the correct interpretation, that is, the one intended by the speaker. We will leave it to the reader to construct an argument showing that, in this situation, the interpretation outlined above is in fact the only interpretation which meets the consistency criterion.

Case 4.G

(26b) - *Coffee would keep me awake* - can of course also be used to convey the implicit information that the speaker wants some coffee. Mary can rationally expect John to interpret the utterance in this way only, however, if the contextual assumption that Mary does want to stay awake is more accessible to him than the assumption that she does not want to stay awake. This will be the case, for instance, if it is mutually manifest to John and Mary that Mary is tired, but needs to finish some work. Such an assumption can become manifest through an explicit utterance of the speaker's at some point prior to the exchange under consideration, or it can be manifest on the basis of the hearer's knowledge that the speaker has an assignment to hand in the following day and that she will have to work all night to finish it in time. In this variant of the situation outlined above, the assumption (31) will therefore be more accessible than the assumption (28).

- (31) *The most readily accessible context*
Mary wants to stay awake.

The first accessible context for the interpretation of (26b) in Case 4.G is then assumption (31), and not (28). The contextual assumption (31) - that Mary wants to stay awake - combines with the propositional form (27) of the utterance in the following deductive argument to yield an implicated conclusion:

- (32) (a) If Mary wants to stay awake, and if drinking some coffee would keep Mary awake, she will want to drink some coffee.
- (b) Mary wants to stay awake.

- (c) Drinking some coffee would keep Mary awake.
- (d) Therefore, Mary wants to drink some coffee.

So the implicature recovered by the hearer in this second case is the following:

(33) *An implicature recovered for utterance (26b)*

Mary wants to drink some coffee.

In the second situation outlined here, the most readily accessible interpretation tested for consistency with the principle of relevance and found to be consistent with it will again be the only such interpretation. Any other intended interpretation will, at the very least, have required some unjustified processing effort for its recovery, thereby failing to meet the least effort presumption.

Two important points about utterance interpretation can now be highlighted with reference to our discussion of Cases 4.F and 4.G.

The first point: the context for the interpretation of a linguistic utterance is not given, but is a variable to be selected by the hearer. The selection is made on the basis of accessibility.⁴² The hearer first attempts to interpret an utterance in the most easily accessible context. If the interpretation yields adequate effects in this context, he will adopt this interpretation as the correct, intended, interpretation. Only if the utterance fails to yield adequate effects in the first accessible context, will the hearer select the next most accessible context, which requires correspondingly more processing effort.

The second point: since interpreting an utterance is for the best part an inferential process, there is no guarantee that the interpretation recovered by the hearer, because consistent with the principle of relevance, is in fact the correct, intended, interpretation. Like all other inferential processes, this one has an element of uncertainty built into it. In the case of implicatures, the indeterminacy of the process of recovery lies specifically in the fact that the speaker has to foresee which contextual assumptions are most accessible to the hearer, and will therefore be used by him to derive contextual effects. Naturally, there is always the possibility that the speaker may make a mistake. Communication may then fail in that the hearer will derive a different implicature from the one intended by the speaker. Suppose, for instance, the speaker by uttering (26b) intends to convey (33) - that she wants to drink some coffee - but is *mistaken* in assuming that the contextual assumption (31) - that she wants to stay awake - is more accessible to the hearer than the assumption (28) - that she does not want to stay awake. The hearer will then derive not (33) but the contrary implicature (30), since it is this latter that will meet the consistency criterion. Clearly, however, in this case the hearer's best hypothesis about the intended interpretation is not the correct interpretation. Communication will have failed. But note that here, too, the failure of communication is accounted for by the consistency criterion.

Communication will also fail in a situation where, contrary to the speaker's expectations, (28) and (31) are equally accessible to the hearer. Being equally accessible, the two interpretations will then yield comparable contextual effects. As a result the hearer will be

⁴² For a brief outline of some of the factors which affect accessibility cf. 4.3.2 above.

unable to choose between the two possible interpretations. Neither interpretation will be consistent with the principle of relevance. In order to determine the speaker's intended interpretation, the hearer will have to ask the speaker explicitly whether she does or does not want to stay awake, as in the following responses to the utterance (26b):

- (34) (a) Well, *do* you or *don't* you want to stay awake?
(b) Well, does that mean that you *do* want coffee or that you *don't* want coffee?

If the speaker could rationally have foreseen that (28) and (31) would be equally accessible to the hearer, she could easily have spared him a great deal of unjustified processing effort by adjusting her utterance. For example, by adding at the start of (26b) a phrase like *Yes, please* or *No, thanks* - depending on what she wanted to convey - the speaker could have enabled the hearer to recover her intended interpretation in return for appreciably less processing effort.

But at this juncture another important question arises: rather than utter *Coffee will keep me awake* and so convey the intended information implicitly, could not the speaker have provided the hearer with this information by answering directly - for example, by saying *Yes, I want some coffee* or *No, I don't want any coffee*? Such a direct answer would have spared the hearer an appreciable amount of processing effort. Would such a direct utterance not, then, be more relevant than the less direct utterance under consideration, given that processing effort co-determines degree of relevance? Sperber and Wilson's (1986:196-7) answer to this is that *the extra processing effort involved in recovering this information via an interpretation of the indirect answer is offset by extra contextual effects*. As they put it: "...it follows from the principle of relevance that the surplus of information given in an indirect answer must achieve some relevance in its own right". What this means in the case of our example is that, for the indirect answer (26b) to the question (26a) to be optimally relevant, the hearer must be able to recover some contextual effects over and above those presented in (30) and (33). If there are no such further contextual effects to be recovered, then the speaker has been neglecting some obviously cheaper way which was open to her for conveying these assumptions - for example, by stating them directly.

Let us briefly consider some further implicatures of (26b) which John could derive. For the purpose of this illustration, we consider only Case 4.F. In that situation, as argued above, the hearer, John, derives the implicature (30) that Mary does not want to drink any coffee. John is also able to recover some further information, however: information about the reason why the speaker does not want any coffee, the reason being that the speaker does not want to be kept awake (rather than - let us say - that the hearer makes awful coffee, or that the speaker does not want to spend more time with the hearer). That is, the indirect utterance conveys not only a refusal of the offer of coffee but also an explanation for this refusal, whereas a direct answer such as a simple *No, thanks* conveys a refusal only. This "extra" information implicitly conveyed by the indirect utterance is, then, what makes up for the extra processing effort needed to interpret the indirect utterance.

But there are other implications which John is able to derive from (26b), if he combines it with some further more or less easily accessible assumptions. In the situation outlined as Case 4.F, given the accessibility of the assumption that tea is also a stimulant, John may

recover the further assumption that Mary does not want tea, either. Similarly for chocolate or any other substance which is itself a stimulant or contains some stimulant.

For Mary to have expected her utterance (26b) to achieve relevance, she must have expected John to recover at least some of these additional assumptions. But note that, because any one of several different implicatures could establish the optimal relevance of Mary's utterance, no single one of these need be specifically intended by Mary. Certainly Mary's answer in (26b) gives her hearer some encouragement to think that she would also decline an offer of tea. But her backing of this implicature is clearly less than her backing of the implicature (30). Could John also derive from (26b) the implication that Mary would like a drink that is not stimulatory, such as herbal tea? If Mary's utterance offers any encouragement at all for deriving this implication, it is certainly much less than that for deriving any of the other implicatures discussed above. Suppose further that John believes that people who avoid coffee for fear of being kept awake also do not do any exercise before going to bed, and that he finds it worth his while to derive the implication that Mary will not exercise before going to bed. It is highly doubtful that Mary has given John any encouragement to derive this implicature. Rather, it seems likely that John will himself have to take all the responsibility for supplying this implicature.

The discussion above, informal and inexhaustive as it may be, serves to illustrate an important point made by Sperber and Wilson (1986:199): the implicatures of an utterance, like assumptions in general, may vary in their strength. They (1986:199) explicate this notion of 'strength of an implicature' as follows:

"To communicate an assumption A is to make mutually manifest one's intention to make A manifest or more manifest. The greater the mutual manifestness of the informative intention to make manifest some particular assumption, the more strongly the assumption is communicated."

These comments must be considered against the background of Sperber and Wilson's (1986:59) view of communication as a matter of degree: assumptions are not seen as either communicated or not communicated; rather, there is a set of assumptions which, as a result of communication, become manifest or more manifest to varying degrees.

The strongest possible implicatures of an utterance are those fully determinate implicatures which must be supplied if the interpretation of the utterance is to meet the consistency criterion. The speaker takes full responsibility for such implicatures. The implicature (30), 'Mary does not want to drink any coffee', is an example of such a fully determinate implicature for which the speaker takes full responsibility. It is fully determinate in the sense that the speaker does not expect her hearer to supply merely something similar to this assumption; rather, she expects him to supply precisely this assumption. The speaker is entirely responsible for its truth, just as much as if she had asserted it explicitly. In the terminology of Sperber and Wilson (1986:59), the assumption is "strongly communicated".

According to Sperber and Wilson (1986:195), not all implicatures of utterances are such fully determinate assumptions for which the speaker is fully responsible. Rather, as in the case of (26b), there may be a further range of implicatures which the speaker's utterance may encourage, but not compel, the hearer to supply. In the case of our example, the

implicatures that the speaker does not want any tea, that the speaker does not want any chocolate, that the speaker does want herbal tea or milk, that the speaker will not exercise before going to bed, fall in this category. The weaker the encouragement given by the speaker, and the wider the range of possibilities among which the hearer can choose, the weaker the implicatures. In the terminology of Sperber and Wilson (1986:59-60): in a case such as this, where the communicator's intention is to increase marginally the manifestness of a wide range of assumptions, each of them is "weakly communicated".

Beyond a certain stage, the speaker is no longer responsible for the inferences made by the hearer. The hearer then accepts full responsibility for supplying them himself. In the case of our example, at least the inference that the speaker will avoid exercising before going to bed falls into this class.

Sperber and Wilson (1986:199) expressly claim that there is no clear cut-off point between wholly determinate, specifically intended inferences and indeterminate, wholly unintended inferences. They claim, in other words, that there is no cut-off point between inferences strongly backed by the speaker and inferences made solely on the hearer's responsibility.

Sperber and Wilson (1986:199) then claim that the indeterminacy of certain implicatures presents no particular formal problem on the relevance-theoretic approach.⁴³ As we have seen, an utterance with a fully determinate implicature forces the hearer to supply just that implicature and to regard it as part of the assumptions which the speaker intends to convey. The recovery of this implicature is forced, of course, by the hearer's attempt to fulfil the presumption that the utterance will be optimally relevant to him. An utterance with a small range of strongly implicated implicatures strongly encourages the hearer to use some subset of these implicatures in recovering an overall interpretation which is optimally relevant to him. An utterance with a wide range of weakly implicated implicatures likewise encourages the hearer to regard some subset of them as part of the interpretation intended by the speaker. But, clearly, the wider the range of possibilities from which the hearer can choose, and thus the weaker the implicatures, the less the confidence the hearer can have that the particular implicatures which he supplies are true to the hearer's thoughts. This is in fact where the indeterminacy lies. The general conclusion that the hearer cannot always be sure that the implicatures recovered by him are true to the speaker's thoughts leads Sperber and Wilson (1986:200) to an important statement about the nature of communication (the italics are ours):

"...people may entertain different thoughts and come to have different beliefs on the basis of the same cognitive environment. *The aim of communication in general is to increase the mutuality of cognitive environments rather than guarantee an impossible duplication of thoughts.*"

That communication serves, not to "duplicate" the speaker's "thoughts" in the hearer's mind, but rather to increase the mutuality of speaker's and hearer's cognitive environments is a view to which we will return immediately below in our discussion of stylistic effects.

⁴³ Cf. their (1986:195ff.) criticism of the assumption generally held in pragmatics that all implicatures are fully determinate.

4.5 How the consistency criterion explains the recovery of stylistic effects

Sperber and Wilson (1986) discuss two types of stylistic effects in some detail. The first type they call *presuppositional effects*. Their claim is that all effects discussed in the literature in terms of pairs of contrasted notions such as 'presupposition' and 'focus', 'topic' and 'comment', 'presupposition' and 'assertion', or 'given' and 'new', can be explained in terms of backgrounding and foregrounding. Backgrounding and foregrounding, in turn, come about as automatic results of the hearer's tendency to maximise relevance and of the speaker's exploitation of that tendency.⁴⁴ The second type of stylistic effects discussed in some detail by Sperber and Wilson (1986) they label *poetic effects*. It is stylistic effects of this second type that will form the focus of our illustration of a relevance-theoretic account of stylistic effects. Before we consider any examples in detail, however, let us consider at a more general level a few important elements of Sperber and Wilson's views on style.

Sperber and Wilson's (1986:202ff.) account of stylistic effects within a relevance-theoretic framework starts from the observation that two utterances with the same linguistically determined truth-conditions need not have identical implicatures. Such a difference in implicatures can arise because utterances with the same linguistically determined truth-conditions may differ in the processing effort they require. From Sperber and Wilson's (1986:202ff.) discussion it is clear that, for them, such differences in processing effort will derive from differences in linguistic structure between the utterances - for example, differences in linear order, differences in focal stress, and the difference between the presence and the absence of a "figure of speech". The principle of relevance, of course, takes processing effort into account; clearly, therefore, the principle implies that differences in required processing effort between two utterances whose linguistically determined truth-conditions are the same may cause them to have different contextual effects. So the general idea, here, is that certain features of its linguistic structure may have an influence on the effort required to process an utterance and, consequently, on the contextual effects produced by the utterance. This idea is the key to an explanatory theory of style, according to Sperber and Wilson (1986:202).

Fundamental to Sperber and Wilson's attempt to develop a relevance-theoretic account of style is their (1986:218) view that no speaker can avoid choosing a style. On this view, moreover, the speaker's choice of style will necessarily reveal her assumptions about the hearer's contextual resources and about the hearer's processing abilities. Sperber and Wilson illustrate this view with the following three utterances:

- (35) (a) Only amateurs can compete in the Olympics.
(b) The Olympic Games is an international sporting competition held every four years. Only amateurs can compete.
(c) The Olympic Games is an international sporting competition held every four years. Only amateurs - that is, people who receive no payment for their sporting activities - can compete in the Olympic Games. Professionals - that is, people

⁴⁴ For an extensive discussion of the role which the principle of relevance can play in accounting for such presuppositional effects cf. S&W 1986:202-217.

who receive some payment for their sporting activities - are not allowed to compete in the Olympic Games.

These utterances, according to Sperber and Wilson (1986:218), do not really differ in their import. Rather, they differ in the amount of help they give the hearer in recovering this import. What the speaker of (35a) trusts the hearer to know about the Olympics is stated explicitly in (35b) and (35c). What the speaker of (35a) and (35b) trusts the hearer to know about amateur status is stated explicitly in (35c). Sperber and Wilson (1986:218) characterise the differences in style among (35a)-(35c) in terms of a notion of 'heaviness': the style of (35c) is heavier than that of (35b), which in turn is heavier than that of (35a). The speaker's choice of one of these styles reveals what degree of mutual understanding she assumes to exist between herself and the hearer. The more the information left implicit, the greater the degree of assumed mutual understanding. A speaker who uses (35a), the least heavy of the three utterances, clearly makes manifest that she assumes a much greater degree of mutual understanding with the hearer than does the speaker who uses (35c), the heaviest of the utterances (Sperber and Wilson 1986:218).

According to Sperber and Wilson, a speaker's assumptions about the degree of mutual understanding between herself and the hearer owe their significance to the principle of relevance; less directly, so too does the speaker's corresponding choice of a more or less explicit form or style for communicating the intended information. Let us see why this is claimed to be so. A speaker aims at optimal relevance. Optimal relevance depends not only on contextual effects, but also on processing effort. The speaker must therefore choose a form which will enable her hearer to recover the intended information with the minimum of processing effort. Specifically, the speaker must decide what to make explicit, and what to leave implicit. Given then that her aim is to achieve optimal relevance, the speaker will leave implicit all information which her hearer can be trusted to supply with less effort than would be required to process an explicit prompt (Sperber and Wilson 1986:218). Conversely, the speaker will communicate explicitly all information of the following two sorts: all information, if any, which she cannot trust her hearer to supply with less effort than would be required by the processing of an explicit prompt, and all information, if any, which she cannot trust her hearer to supply at all. So, by making the right assumptions about the degree of mutual understanding between herself and the hearer, and by choosing a linguistic structure whose form strikes the right balance between explicit and implicit information, the speaker ensures that the hearer is able to recover the intended interpretation in return for the minimum of processing effort. To put it another way: the speaker, in so doing, ensures that the utterance achieves optimal relevance.

To prevent any misunderstanding of what is a crucial point, let us emphasise it here: there is no guarantee that the speaker will estimate correctly the degree of mutual understanding between herself and the hearer. She may overestimate the degree of mutual understanding, so making it difficult or even impossible for the hearer to recover her intended interpretation. Conversely, she may underestimate the degree of mutual understanding. In such a case her explicit prompts, although intended to be helpful, may seem patronising or even offensive to the hearer (Sperber and Wilson 1986:218). The crucial point here is not about correctness, however, but about inevitability: the speaker cannot avoid making assumptions about the degree of mutual understanding between herself and the hearer, nor can she avoid choosing a form that reveals these assumptions. To quote Sperber and

Wilson's (1986:218) succinct and memorable phrasing of this crucial point: "There is no entirely neutral style".

As well as differing in their degree of reliance on implicature, styles, according to Sperber and Wilson (1986:218), can differ in the degree to which they constrain or guide the hearer's search for relevance. This latter kind of difference can be briefly illustrated by comparing a direct with an indirect answer to a question.⁴⁵

- (36) (a) *Peter*: Would you drive a Mercedes?
(b) *Mary*: I wouldn't drive a Mercedes.
(c) *Mary*: I wouldn't drive any expensive car.

A direct answer to question (36a), such as (36b), leaves the hearer free to process the information contained in the propositional form in whatever way he likes: the speaker does not give him any guidance on what further implications he should draw from this information. In contrast, the indirect answer in (36c) suggests a particular line of processing for the hearer's computation of contextual effects. In saying (36c), for instance, Mary not only expects Peter to access and use the assumption that a Mercedes is an expensive car. She also encourages him to speculate on, and to derive some additional conclusions from, the assumption that she will not drive any expensive car. For instance, it will be reasonable for Peter to assume that Mary has also intended to convey to him that she will not drive a Rolls Royce or a Cadillac, given the accessibility of the stereotyped item of general knowledge that a Rolls Royce and a Cadillac are very expensive. In sum: the indirect answer (36c), in contrast to the direct answer (36b), guides and constrains the hearer's search for relevance.

Sperber and Wilson's (1986:219) view of the relation between style and relevance is neatly captured in the following formulation: "Style arises...in the pursuit of relevance". One of the classical figures of style which they use to make this view concrete is that of repetition (also called *epizeuxis*). Let us take up now their relevance-theoretic account of repetition. Doing so will give us an opportunity to present in some detail their account of poetic effects.

Consider, then, the task of a hearer faced with an utterance which contains a repetition of some expression. Every utterance, let us recall, carries a guarantee of optimal relevance. So the hearer must somehow reconcile the fact that a certain expression has been repeated with the assumption that optimal relevance has been aimed at (Sperber and Wilson 1986:220). A repeated expression requires extra processing effort on the hearer's part. In order to ensure optimal relevance, this extra processing effort must be offset by an increase in contextual effects triggered by the repetition itself (Sperber and Wilson 1986:220). The hearer faced with the task of interpreting an utterance which contains a repetition must therefore determine what effects, in addition to those communicated by the first instance of the repeated element, the speaker could rationally have expected him to derive. Since such an increase in effects can be achieved in various ways, repetitions can have a variety of effects on utterance interpretation. Let us review briefly some of these effects.

⁴⁵ For a much more extensive discussion of this example cf. S&W 1986:194ff. The same points can be made with reference to the "coffee" example discussed in 4.4 above.

For example, what interpretation of the repeated adjective in (37b) would be consistent with the principle of relevance?

- (37) (a) We went for a long walk.
(b) We went for a long, long walk.

(37b) would normally be taken to convey that the speaker and her companion(s) went for a very long walk. In this case, it would be consistent with the principle of relevance for the hearer to assume that what the speaker intended the repetition to communicate was the following: the walk was longer than the hearer, in the absence of the repetition, would have thought. Here, clearly, the repetition contributes to the propositional form or, equivalently, the explicit content of the utterance (Sperber and Wilson 1986:219). The main point, then, is this: the extra effort needed to process the second occurrence of *long* is offset by the contribution which the repetition of the adjective makes to the utterance's explicit content.

Now compare (37b) with (38b).

- (38) (a) I shall never smoke again.
(b) I shall never, never smoke again.

Here it would be consistent with the principle of relevance for the hearer to assume that the speaker attached a higher confirmation value to the assumption expressed than the hearer would otherwise have thought (Sperber and Wilson 1986:220-221). The speaker, aware that her utterance will meet with scepticism, repeats the word *never*, the scepticism being liable to focus on precisely the absoluteness of the meaning of this word. The speaker makes use of this repetition in an attempt to convince the hearer that she means what she says. So *never, never* is similar in import here to *definitely never*, and "reflects the speaker's degree of commitment to the assumption expressed" (Sperber and Wilson 1986:220). The repetition of the adverb therefore *strengthens* the explicit content, and consequently all the contextual implications, of the utterance. Here, then, is the main point in the case of (38b): the extra processing cost incurred by the second occurrence of *never* is offset by the strengthening of the utterance's explicit content and, hence, the strengthening of all its contextual implications.

In some cases the increase in processing effort required by a repetition in an utterance - unlike (37b) - cannot be offset by enriching the propositional form of the utterance or - unlike (38b) - cannot be offset by strengthening the implicatures of the utterance. Consider for instance the effect of the repetition of the noun *fox* in (39b).

- (39) (a) There's a fox in the garden.
(b) There's a fox, a fox in the garden.

Here the repetition cannot be accounted for by assuming that there is more than one fox in the garden, or by strengthening the assumption that there is a fox in the garden. Sperber and Wilson (1986:221) suggest that in such cases the repetition yields an increase in contextual effects by encouraging the hearer to expand the context in which he is to process the utterance, thereby deriving further implicatures. In this particular case, Sperber and Wilson (1986:221) suggest furthermore, the hearer "is being encouraged to dig deeper into his

encyclopaedic entry for *fox*, with a guarantee that the extra processing effort will be outweighed by a gain in contextual effects: the fact that there is a fox in the garden is presented as more relevant than the hearer would spontaneously have realised". For instance, the hearer, by paying attention to the fact that there is fox in the garden, by making the effort to call to mind certain basic facts about foxes, and by processing the readily accessible information that there are chickens in the garden is likely to derive some strong and predictable contextual implications, such as "The chickens are in danger." Such implications are likely to be interpreted as strong implicatures of the utterance in the sense set out towards the end of 4.4. That is, they are highly determinate assumptions for which the speaker takes a large share of the responsibility.

We can now return to the repetition used in 2.4.3 to illustrate the notion of a 'poetic effect'. For convenience, the example - first presented as (8) in 2.4.3 above - is repeated here:

- (40) (a) My childhood days are gone.
(b) My childhood days are gone, gone.

In the case of (40b), as in that of (39b), the repetition cannot be accounted for by enriching the propositional form of the utterance or by strengthening its implicatures. To be more specific: the repetition cannot be accounted for by assuming that the speaker's childhood days are longer gone, or more definitely gone, than might otherwise be assumed. As in the case of (39b), so here the repetition serves as an encouragement to the hearer to expand the context, and to derive, by so doing, additional implications which will compensate for the additional processing effort. But here, in contrast to (39b), the utterance does not make a specific assumption strongly manifest to the hearer in an expanded context. Rather, the repetition here serves to increase marginally the manifestness of a great many weakly manifest assumptions. For instance, the repetition may encourage the hearer to compare the speaker's childhood and her present condition, to assume that the speaker is reminiscing and making a similar comparison, and to imagine the feelings this may evoke in her. In this case, unlike (39b), the utterance does not achieve its relevance through a strongly manifest additional implicature. Rather, it achieves its relevance through a wide array of weak implicatures. Although the speaker gives some backing to the contextual effects, the hearer takes a large share of responsibility in imagining what it may be like for the speaker to be way past her youth. In the case of (40b), then, we have another instance of weak communication in the sense defined by Sperber and Wilson (1986:59-60).

Sperber and Wilson (1986:222) use the term *poetic effect* to refer to the peculiar effect of utterances, such as (40b), which achieve most of their relevance through a wide array of weak implicatures. How do such poetic effects modify the mutual cognitive environment of speaker and hearer? To this question Sperber and Wilson (1986:224) respond as follows:

"[Poetic effects] do not add entirely new assumptions which are strongly manifest in this environment. Instead, they marginally increase the manifestness of a great many weakly manifest assumptions. In other words, poetic effects create common impressions rather than common knowledge. Utterances with

poetic effects can be used precisely to create this sense of apparently affective rather than cognitive mutuality. What we are suggesting is that, if you look at these affective effects through the microscope of relevance theory, you see a wide array of minute cognitive effects."

Clearly, this answer turns on their (1986:200) view of the nature of communication, quoted at the end of 4.4: communication serves, not to "duplicate" the speaker's "thoughts" in the hearer's mind, but rather to increase the mutuality of speaker's and hearer's cognitive environments.

In sum, then: Sperber and Wilson (1986:217-224) claim that the criterion of consistency with the principle of relevance is able to account for stylistic effects in general. Stylistic differences, according to them (1986:224), "are just differences in the way relevance is achieved". Moreover, Sperber and Wilson (1986:224) claim that the criterion of consistency with the principle of relevance can also account for the interpretation of utterances which are meant to achieve poetic effects. This is possible, given Sperber and Wilson's (1986:59-60) notion of 'weak communication': weak communication is communication in which the manifestness of a large range of assumptions is only marginally increased.

4.6 How the consistency criterion explains the recovery of a metaphorical interpretation

The hearer's task - recovering the set of assumptions $\{I\}$ which the speaker intends to make manifest to him by means of her utterance - includes, as we illustrated in 2.4.4, the subtask of recovering any intended metaphorical interpretation. Recovering an intended metaphorical interpretation requires, so Sperber and Wilson (1986:231ff.) argue, no special interpretive abilities or procedures on the hearer's part: an intended metaphorical interpretation, like other parts of the intended interpretation of an utterance, is recovered as a result of the hearer's search for optimal relevance. In this section, then, we will present and illustrate Sperber and Wilson's view that the recovery of intended metaphorical interpretations is explained by the consistency criterion.

Underlying Sperber and Wilson's relevance-theoretic account of metaphor is their conception of the relationship between the propositional form of an utterance and the thought of the speaker which this utterance represents. To start off, therefore, let us try to gain a clear idea of this conception.

That an utterance can represent some state of affairs in virtue of its propositional form's being true of that state of affairs is of course a familiar enough notion. In Sperber and Wilson's (1986:228) terminology, the representation in this case is a *description*, or is used *descriptively*. Recall, for instance, utterance (16) and the circumstances of Case 4.B. We can now note three points about them, making use of the terminology just mentioned. Firstly, in the circumstances of Case 4.B, utterance (16) represents a state of affairs - namely, the state of affairs that an individual named Peter is batting with a cricket bat which is too big for him to bat with successfully. Secondly, utterance (16) is enabled to represent this state of affairs in virtue of its propositional form's being true of it. Thirdly, the representing which utterance (16) does in this way is a description of the state of affairs

represented; or, equivalently, the representation which utterance (16) provides in this way is used descriptively of the state of affairs represented.

According to Sperber and Wilson, however, an utterance which represents a thought of a speaker's does not do so by being a description of this thought. Rather, an utterance represents a thought in virtue of its resembling this thought in certain respects. Representation by resemblance is used quite often in non-verbal inferential communication. For instance, Mary may invite Peter for a drink by imitating the act of drinking. To understand Mary's intention fully, Peter has to notice a similarity between Mary's action and the act of drinking. Mary's action and the act of drinking resemble each other to the extent that they have properties in common, and it is by exploiting these resemblances that Mary's action represents the act of drinking.

In the case of verbal inferential communication, utterances can represent phenomena by the exploitation of linguistic resemblances. Consider for instance Mary's utterance (41b).

- (41) (a) *Peter*: What is the last line of 'Rule Britannia'?
(b) *Mary*: 'Britons never never never shall be slaves.'

Here Mary is clearly not expressing some belief of her own; rather, she is merely reproducing the words of a song. In short, Mary's utterance is a direct quotation. Her utterance represents the line of the song in virtue of its resemblance to the form of that line.

But an utterance not only has form; it also has content. More precisely, it has a propositional form. Another way therefore in which an utterance is able to represent some other representation, provided this latter also has a propositional form, is by means of a resemblance between the two propositional forms. If a representation with a propositional form, such as an utterance, represents some other representation with a propositional form in virtue of a resemblance between those propositional forms, the first representation is said to be an *interpretation* of the second one, or to be used *interpretively* (Sperber and Wilson (1986:228-229)).

A thought, according to Sperber and Wilson (1986:229), is an instance of a representation which has a propositional form. Any utterance which represents a thought of a speaker, then, is in their view a public interpretation of this thought. That is, an utterance represents a speaker's thought in virtue of a resemblance between the propositional form of the utterance and the propositional form of the thought. Sperber and Wilson (1986:230-231) argue, moreover, *that every utterance is an interpretive expression of a thought of the speaker's*: every utterance, that is, represents a thought of the speaker's in virtue of a resemblance between the propositional form of the utterance and the propositional form of the thought.

What does it mean for the propositional form of an utterance to resemble that of a thought? Sperber and Wilson (1986: 228;233;234) answer this question as follows: the propositional form of an utterance resembles that of a thought to the extent *that they share logical properties, specifically, logical and contextual implications.*

How closely must the propositional form of an utterance resemble that of the thought which it interprets? Sperber and Wilson (1986:229) explicitly deny that identity between the two propositional forms is necessary. Rather, they claim that the degree of resemblance varies from case to case, but is always determined by the principle of relevance. How, will be explained directly below.

Sperber and Wilson (1986:233ff.) hold that an utterance, in its role as an interpretive expression of a speaker's thought, is strictly literal if it has the same propositional form as this thought. That is, in the case of a literal utterance, the set of implications $\{I\}$ of the speaker's thought will be exactly identical to the explicitly and contextually determined content of the utterance. But, according to Sperber and Wilson (1986:232), such identity of propositional forms, or literality, is the limiting case rather than some norm.

To say that an utterance is less than strictly literal is to say that its propositional form shares some, but not all, of its logical properties with the propositional form of the thought which it is being used to represent. Sperber and Wilson (1986:233) claim that, from the standpoint of relevance theory, there is no reason to assume that the optimally relevant interpretive expression of a thought is always the most literal one. Rather, they claim, there are many quite ordinary situations in which a literal utterance is not optimally relevant. Consider for instance the alternative answers (42b) and (42c) to the question (42a).⁴⁶

- (42) (a) *Peter: How far is Nottingham from London?*
(b) *Mary: 120 miles.*
(c) *Mary: 118 miles.*

If Mary knows that Nottingham is exactly 118 miles from London, then she has a choice between the strictly literal and truthful answer in (42c) and the less literal, strictly false, answer in (42b). Suppose that Mary knows that Peter always drives at 60 miles an hour, and that he is trying to decide at what time to leave London for dinner in Nottingham. The answers in (42b) and (42c) will both yield the contextual implication that driving from London to Nottingham would take him about two hours, and that he should plan his journey accordingly. Processing (42c) will require more mental effort, however, than processing (42b): mental calculations in round figures require less effort. When the effort side of the principle of relevance is brought into the picture, the less than literal answer (42b) therefore turns out to be more relevant than the literal answer (42c). The two answers yield the same contextual implication, but (42b) does so with less effort.

In general, Sperber and Wilson (1986:233ff.) argue, a hearer should take an utterance as strictly literal only if nothing less than full literality confirms the presumption of optimal relevance. In general, some non-literality or looseness of expression is to be expected. For instance, we do not take someone to task for giving us the time as 5 p.m., if in fact it is two minutes before or after 5 p.m., unless the relevance of the utterance depends on such exactitude.

Given then that all that the hearer can take for granted is that the speaker intends her utterance as an interpretation (or, equivalently, an interpretive expression) of one of her

⁴⁶ For this illustration of loose use cf. W&S 1986:82 .

thoughts, how is the hearer to work out which subset of the utterance's logical and contextual implications are implications of precisely that particular one of her thoughts which the speaker intends to convey to him? Here, in full, is Sperber and Wilson's (1986:234) answer to this question:

"If the speaker has done her job correctly, all the hearer has to do is start computing, in order of accessibility, those implications which might be relevant to him, and continue to add them to the overall interpretation of the utterance until it is relevant enough to be consistent with the principle of relevance."

Let us translate this answer into the terminology of 4.3 above - that is, into the terms of the task-specific criterion (11). Then what this answer says is that the correct task-specific interpretation is the first accessible interpretation which leads to an overall interpretation which meets the consistency criterion (10). In order to recover this interpretation, the hearer must select the most accessible assumption, and test whether it yields an overall interpretation which meets the consistency criterion (10). If it does, the hearer must assume that the interpretation selected and tested is the correct, intended, interpretation. If it does not, he should select and test the next most accessible assumption. And the hearer is to repeat this process of selection and evaluation until he has constructed an overall interpretation which meets the consistency criterion. In sum: recovering implicatures in cases of the loose use of language - that is, any less than fully literal use of language - is in all essential respects similar to the process of recovering the implicatures of an ordinary, literal, assertion - the process of recovery which we outlined in 4.4.

Let us return now to the case of metaphorical interpretation. Sperber and Wilson (1986:235) claim that there is no discontinuity between the loose use of language discussed above and a variety of "figurative" uses, including metaphor. In all these cases, the propositional form of the utterance differs to a greater or to a lesser degree from that of the thought of which it is an interpretation (or, equivalently, an interpretive expression). That is, the implications of the utterance are not identical to those of the thought. In both cases, however, the hearer can proceed on the assumption that the two propositional forms have some identifiable logical and contextual implications in common. And in both cases these implications will be recovered along the lines sketched above. The two cases - namely, the loose use of language and figurative use such as metaphorical use - therefore involve the same interpretive abilities and procedures.

Let us see now how this works in the case of utterance (43), first presented as (9) in 2.4.4:

(43) This room is a pigsty.

According to Sperber and Wilson (1986:236), such instances typically give access to an encyclopaedic schema with one or two dominant and, hence, most highly accessible assumptions. In terms of the most highly accessible assumptions in the encyclopaedic entry for the concept of 'pigsty', for example, pigsties are stereotyped as being extremely dirty and untidy. When processed in the context of such assumptions, (43) yields the contextual implication that the room concerned is extremely dirty and untidy. Precisely this conceptual implication will be the first accessible task-specific interpretation of (43) to be evaluated in terms of the consistency criterion. Accordingly, if the speaker did not intend this contextual

implication to be derived, she should have picked an utterance which would not give rise to this implication. She should have said something like the following, for example:

(44) This room is where the pigs are kept.

But note that, in processing a metaphorical expression of the assumption that the room is extremely dirty and untidy, the hearer needs to expend more cognitive effort than in processing a literal expression of this, such as the following:

(45) This room is extremely dirty and untidy.

In cases such as (43), of course, the principle of relevance gives rise to a question: what justifies the extra processing effort required? In terms of the principle, one would expect the extra processing effort to be offset by some extra gain in contextual effects. In the case of our example (43), the metaphorical expression conveys the idea that the room in question is dirty and untidy beyond the norm; that is, beyond what the speaker could have satisfactorily conveyed by saying merely that the room was extremely dirty and untidy. In order to recover an overall interpretation which a rational speaker might have expected to be optimally relevant to the hearer, the hearer will therefore have to add the assumption that the room is dirty and untidy beyond the norm to the set of assumptions which he takes as the overall interpretation of the utterance.

In the case of the example discussed above, the contextual effect takes the form of a strong implicature; this is due to the highly stereotyped nature of the assumptions involved. But the same kind of account applies to less stereotyped, more creative, metaphors. Consider, for instance, the following utterance, also discussed by Sperber and Wilson (1986:236):

(46) Robert is a bulldozer.

Assume that the entity referred to in this utterance as *Robert* is identifiable as a person. Then the metaphor in this utterance brings together two encyclopaedic entries which are not normally combined in a subject-predicate relationship. Many of the contextual implications yielded by this combination will be contradictory and, as a result, will be discarded immediately. For instance, the implication - which follows from the meaning of *bulldozer* - that Robert is made of metal contradicts the implication - which follows from the reference assigned here to *Robert* - that Robert is a human being. In contrast with (43), namely *This room is a pigsty*, (46) does not make a single, strong implicature available as the intended interpretation. Rather, (46) brings to mind a slightly weaker, less determinate, range of implicatures - implicatures having to do with Robert's persistence, his obstinacy, his insensitivity and his refusal to be deflected from his chosen path. In this case the hearer must bear a greater responsibility for the resulting interpretation than he does in the case of (43). According to Sperber and Wilson (1986:236), a good creative metaphor is precisely a metaphor in the interpretation of which a variety of contextual effects can be retained by the hearer and understood by him as weakly implicated by the speaker. The crucial point here is that, in the case of metaphorical utterances, once again the consistency criterion proves its explanatory power: the consistency criterion explains how it is possible for the hearer to recover the intended metaphorical interpretation of creative metaphors such as that in (46).

As this exposition of their account of the metaphorical interpretation of (46) is intended to make clear, Sperber and Wilson, in explaining the interpretation of the full range of metaphors, again appeal to their notion of 'weak communication' (a notion which we presented in 4.4 above). As they (1986:60) put it, with certain weaker forms of communication "the communicator can merely expect to stir the thoughts of the audience in a certain direction". This, then, according to Sperber and Wilson, is what happens in the case of a creative metaphor such as (46). Their (1986:236-237) comments on the interpretation of a "good creative metaphor" shed valuable light on their ideas about this issue, and we therefore quote these comments at some length:

"A good creative metaphor is precisely one in which a variety of contextual effects can be retained and understood as weakly implicated by the speaker. In the richest and most successful cases, the hearer or reader can go beyond just exploring the immediate context and the entries for concepts involved in it, accessing a wide area of knowledge, adding metaphors of his own as interpretations of possible developments he is not ready to go into, and getting more and more very weak implicatures, with suggestions for still further processing. The result is a quite complex picture, for which the hearer has to take a large part of the responsibility, but the discovery of which has been triggered by the writer. The surprise or beauty of a successful creative metaphor lies in this condensation, in the fact that a single expression which has itself been loosely used will determine a very wide range of acceptable weak implicatures."

In sum: according to Sperber and Wilson the interpretation of metaphor requires no special interpretive abilities on the hearer's part. Metaphor (along with other figurative uses of language) falls on a continuum with quite ordinary cases of the loose (that is, the less than fully literal) use of language, and in its turn the loose use of language is not essentially different from the fully literal use of language. The hearer's recovery of an intended metaphorical interpretation is simply the result of his search for an overall interpretation which a rational speaker might have expected to be optimally relevant to him.

4.7 How the consistency criterion explains the recovery of an ironical interpretation

As was illustrated in 3.3.4 above, a further part of the hearer's task in interpreting an utterance is to recover whatever ironical interpretation the speaker intends. As in the case of metaphor, so in the case of irony Sperber and Wilson (1986:237ff.) claim that relevance theory is able to provide an explanatory account without *any* special theoretical mechanisms over and above the consistency criterion. As in the case of metaphorical utterances, so in the case of ironic utterances they argue that these fall together with some range of quite ordinary utterances - in the present case, the class of echoic utterances.

Sperber and Wilson (1986:230-231) argue, we noted in 4.6, that every linguistic utterance is an interpretive expression of a thought of the speaker's. But what does this thought represent? In the case of ordinary assertions, this thought represents an actual state of affairs - a possibility which we sketchily illustrated at the start of 4.6 above. There are other possibilities, too, however. There is the possibility, for instance, that the speaker's thought

may itself be an interpretation of some other thought. This other thought may be either a thought of someone else's, or a thought of the speaker's own from the past.⁴⁷

Reported speech is a clear case where an utterance is an interpretation of a thought of the speaker's, a thought which is itself an interpretation of some other thought. In such a case the utterance achieves relevance by informing the hearer of the fact that so-and-so has said something, or thinks something. But there are also cases where the utterance achieves relevance by informing the hearer that the speaker not only has in mind something said or thought, but also adopts a certain attitude to it. All utterances which achieve relevance in this way are referred to by Sperber and Wilson as *echoic utterances*. Let us glance at a situation, taken over from Sperber and Wilson (1986:238), in which a simple example of an echoic utterance is produced by a speaker called *Mary*:

- (47) (a) *Peter*: The Joneses aren't coming to the party.
(b) *Mary*: They aren't coming, hum. If that's true, we might invite the Smiths.

In her utterance of the sentence *They aren't coming, hum*, Mary echoes what Peter has just said. That is, Mary's first sentence-utterance is an interpretation of a thought previously expressed by someone other than the speaker. In this case, Mary's first sentence-utterance echoes a thought previously expressed by the hearer, Peter. Clearly, however, this utterance of Mary's does not achieve its relevance simply by reporting what Peter has just said. There is obviously no need, for instance, for Mary to remind Peter of what he has just said. Rather, Mary's first sentence-utterance achieves relevance by conveying to Peter the information that she has been paying attention to his utterance and is weighing up its reliability and implications. Or, to put the point in slightly more general terms, this utterance of Mary's provides evidence of her attitude towards the thought echoed in this utterance. It is in virtue of communicating to Peter this attitude of hers that her utterance achieves relevance.

Verbal irony, in terms of Sperber and Wilson's (1986:239) analysis, is a subcase of echoic interpretation. For them, verbal irony always involves an echoic allusion to a thought held by someone else (or by the speaker herself at some time in the past). And verbal irony, for them, specifically involves the implicit expression by the speaker of a dismissive or disapproving attitude towards the opinion echoed. There are various clues which the hearer may use as evidence when trying to determine the speaker's attitude to the thought or opinion expressed. These clues include tone of voice, paralinguistic features, and context. An ironical utterance, therefore, owes its relevance (at least in part) to the information conveyed by the utterance about the speaker's attitude towards the opinion echoed in the utterance.

Let us consider how this analysis of Sperber and Wilson's applies to a non-ironical utterance, and a corresponding ironical utterance, of the sentence *It's a lovely day for a picnic, indeed*. The ironical utterance was first presented as (10) in 2.4.5 above. Naturally enough, the two utterances are embedded in different situations:

⁴⁷ For a schema systematising all the possibilities which S&W distinguish cf. S&W 1986:232. In the next section we will consider a few more of the possibilities represented there.

- (48) (a) *Peter*: It's a lovely day for a picnic.
[They go for a picnic and the sun shines.]
(b) *Mary* (happily): It's a lovely day for a picnic, indeed.
- (49) (a) *Peter*: It's a lovely day for picnic.
[They go for a picnic and it rains.]
(b) *Mary* (sarcastically): It's a lovely day for a picnic, indeed.

Nevertheless, although these two utterances of the *indeed* sentence are clearly distinct, they also have at least two features in common. Firstly, in both utterances, Mary is echoing an opinion expressed earlier by Peter. Secondly, both utterances provide Peter with some evidence about Mary's attitude towards the opinion echoed. This evidence includes Mary's different tones of voice and the easily accessible, but different, contextual assumptions about the current state of the weather. In situation (48), these factors provide evidence that Mary is in agreement with the opinion she echoes. In situation (49), however, the evidence from these factors supports the inference that Mary not merely disagrees with the opinion she echoes, but in fact rejects it with scorn.

In the three situations considered so far in this section, as will have been obvious, the potential/conceivable interpretations always include one which is more readily accessible than any of the others. (Simply put, as we noted in 4.2, this is the interpretation which is the first to occur to the hearer's mind.) For instance, in situation (48), given Mary's happy tone of voice and given the observable information that the weather is fine, the most readily accessible of the potential interpretations is that Mary is expressing agreement with the opinion echoed. Similarly, in situation (49), Mary's sarcastic tone of voice, coupled with the fact that it is raining, makes it the most readily accessible of the conceivable interpretations that she is expressing disagreement with the opinion echoed. In short: of all the conceivable interpretations, the one most readily conceived by the hearer's mind is the first one to be tested by his mind for consistency with the principle of relevance. And so far, so good, because this is exactly as is required by the consistency criterion (10). If, moreover, this first-tested interpretation yields adequate contextual effects, then of course the consistency criterion (10) requires the hearer to take it that this is the intended interpretation. For instance, once Peter in (49) has established that Mary is echoing and rejecting his own opinion, he can infer that she intends him to recover implicatures such as the following: that he was mistaken in describing the day as a lovely one for a picnic, that they should never have set out, that it is his fault that their day has been ruined, and so on. Peter's recovery of adequate contextual effects such as these implicatures will then have made it worth his while to have processed Mary's utterance along the lines indicated just above.

Our illustration of how the consistency criterion guides the recovery of the implicatures of the ironical utterance in (49b) highlights what Sperber and Wilson (1986:240) call the three "common factors" in the interpretation of ironical utterances:

"The recovery of these implicatures depends, first, on a recognition of the utterance as echoic; second, on an identification of the source of the opinion

echoed; and third, on a recognition that the speaker's attitude to the opinion echoed is one of rejection or disapproval."

With all the echoic utterances presented above, the opinion being echoed is a precisely attributable thought. According to Sperber and Wilson (1986:238), however, this is not a necessary property of echoic utterances. Echoic utterances - including ironical utterances - can interpret the thought of a certain type of person, or of people in general. Suppose, for instance, that the speaker and hearer are watching a game of tennis, which is eventually won by a player not popular with the spectators. As the winner leaves the court, the spectators vocally express their displeasure. The speaker then produces the following utterance:

(50) Everyone loves a winner.

In this situation, (50) is clearly intended ironically. But note that the opinion being echoed here is not attributable to a specific source. Rather, it is an interpretation of a traditional piece of wisdom, attributable to people in general.⁴⁸

One of the principal conclusions which Sperber and Wilson (1986:242) draw from their analysis of irony is that ironical utterances, far from forming a special class of utterances, in fact fall on a continuum with other echoic utterances in regard to the type of attitude expressed. Sperber and Wilson (1986:240) argue that it is futile to try to determine exactly the range of dismissive or disapproving attitudes conveyed by verbal irony. Indeed, they doubt that there is a well-defined subset of ironical attitudes or a well-defined subset of ironical utterances which express these attitudes. From a pragmatic point of view, what is important is that a speaker can use an echoic utterance to convey a whole range of attitudes or even emotions, varying from outright acceptance or approval to outright rejection or disapproval. The hearer's recognition of such attitudes or emotions may be crucial to the interpretation process (Sperber and Wilson 1986:240).

Sperber and Wilson (1986:240) argue also that their account of irony, as outlined above, is superior to the classical account of irony as saying one thing and meaning another. For instance, the classical account cannot explain why a speaker who, in principle, is able to express her intended message directly should decide instead to say the opposite of what she means. On the relevance account, however, the main point of an ironical utterance is clear: to express an attitude of ridicule/rejection/disapproval towards some opinion echoed in the utterance.

But there is another problem with the classical account of irony not faced by the relevance account: many cases of irony simply do not fall within the scope of the classical definition. True enough, however, examples (49b) and (50) above of ironical utterances do happen to fit in with the classical characterisation of irony as saying one thing and meaning the opposite. So, to round off this account of irony within the relevance framework, let us examine briefly an example of an ironical utterance which falls outside the scope of the classical definition, but which the relevance account is able to explain.⁴⁹

⁴⁸ For a similar example cf. S&W 1986:238-239.

⁴⁹ This example is discussed by S&W (1986:241-242).

- (51) When all was over and the rival kings were celebrating their victory with Te Deums in their respective camps....

On the classical account of irony - of saying one thing and meaning the opposite - (51) would have to be treated as equivalent to (52a) or (52b).

- (52) (a) When all was over the rival kings were not celebrating their victory with Te Deums in their respective camps....
(b) When all was over and the rival kings were bewailing their defeat with Misereres in their respective camps... (Voltaire, *Candide*)

According to Sperber and Wilson (1986:241), Voltaire is not suggesting that neither side won the battle and celebrated victory, nor that both sides lost the battle and bewailed their defeat. What Voltaire is doing is to echo claims made by the rival kings. Since these claims contradict each other, he must believe, and must expect his audience to believe, that at least one of these claims is false. Having recognised that Voltaire is echoing an opinion invariably held by both sides after a battle, the hearer/reader must determine what Voltaire's attitude towards this opinion is. Clearly, in this case it is one of scorn and cynicism. In short: according to Sperber and Wilson relevance theory is able to explain ironical utterances which fall outside the scope of the classical characterisation of irony.

Like metaphor then, so Sperber and Wilson claim, irony requires no special interpretive abilities on the part of the hearer. Ironical utterances, like metaphorical utterances, fall on a continuum with a range of quite ordinary utterances - in this case, echoic utterances. Sperber and Wilson (1986:243) draw two interesting conclusions from their account of irony and metaphor. To conclude this section, we quote their comments without further discussion:

"...first, metaphor and irony are not essentially different from other types of 'non-figurative' utterances; and second, they are not essentially similar to one another. Metaphor plays on the relationship between the propositional form of an utterance and the speaker's thought; irony plays on the relationship between the speaker's thought and a thought of someone other than the speaker. This suggests that the notion of a trope, which covers metaphor and irony and radically distinguishes them from 'non-figurative' utterances, should be abandoned altogether: it groups together phenomena which are not closely related and fails to group together phenomena which are."

4.8 How the consistency criterion explains the recovery of the illocutionary force of utterances

We come now to the last subtask of utterance interpretation distinguished in chapter 2: recovering the intended illocutionary force of an utterance. As in the case of all the other tasks distinguished there, Sperber and Wilson (1986:248ff.) claim that no special theoretical mechanisms are required to explain the hearer's successful performance of this task. What

guides the hearer in recovering the intended illocutionary force of an utterance, according to them, is the familiar criterion of consistency with the principle of relevance.⁵⁰

An important element of Sperber and Wilson's relevance-theoretic account of speech acts is their (1986:246) claim that only a small class of speech acts are "of genuine interest to pragmatics" - including the acts of *saying*, *telling*, and *asking*. In making this claim, they reject one of the basic assumptions of speech-act theory, namely that speech acts are, in the words of (Levinson 1983:226), "one of the central phenomena that any general pragmatic theory must account for". They (1986:243) formulate their own position as follows:

"The vast range of data that speech-act theorists have been concerned with is of no special interest to pragmatics. What is of interest is their attempt to deal with the interpretation of non-declarative (e.g. interrogative and imperative) sentences, which must indeed be accounted for in any complete pragmatic theory."

Sperber and Wilson's arguments against the basic assumption referred to above, and in favour of their own position, are based on considerations which are quite independent of relevance theory. For this reason, we will not attempt to present these arguments here. The interested reader may wish to consult Sperber and Wilson's (1986:243-246) own presentation. In the rest of this section, we focus on two members of the small class of speech acts which in their opinion are of genuine interest to pragmatics, namely *saying* and *telling*. We will take up first their sketch of their own proposal for the analysis of *saying*. In presenting this, we will find it convenient to refer back quite extensively to the various examples of utterance interpretation presented in 4.3-4.7 above.

Imagine that a speaker, by uttering a certain sentence, has said that *P*, where *P* is the propositional form of the sentence. Included in the set of assumptions {*I*} which the hearer then has to recover will be the following:

(53) *An assumption scheme*

The speaker has said that *P*.

Sperber and Wilson (1986:246) refer to (53) as an *assumption schema*, to be completed in the case of a given utterance by replacing the variable *P* with the recovered propositional form of the utterance. Sperber and Wilson 1986:(247) suggest that *saying that P* (where *P* is the propositional form of the utterance) must be defined as "communicating that the thought interpreted by *P* is entertained as a description of an actual state of affairs". This definition relies on two notions presented in earlier parts of our text. The first is the notion that an utterance or, more precisely, its propositional form is an interpretation of a thought of the speaker's. This notion was presented in 4.5 above. The second notion is that a thought may represent an actual state of affairs. This notion was presented in 4.6. For convenience, we present the essence of Sperber and Wilson's definition of *saying that P* as follows:

⁵⁰ S&W (1986:254) characterise their discussion of speech acts as "very sketchy"; our attempt to present and illustrate their ideas will, of necessity, reflect the sketchiness of our source.

(54) *Definition of saying that P*

To say that *P* is to communicate that the thought interpreted by *P* is entertained as a description of an actual state of affairs.

As will become clear from the discussion below, all the utterances whose interpretation is discussed in 4.3-4.7 are instances of *saying that P*, where *P* is the recovered propositional form. That is, in all the instances of utterance interpretation presented in 4.3-4.7 the speaker intends to communicate that the thought interpreted by *P* is entertained as a description of an actual state of affairs.

Let us start off by looking at one of the most straightforward cases, namely the utterance of (16) - *Peter's bat is too big* - in the situation outlined in Case 4.B above. In this situation, the speaker clearly intends the hearer to recover the assumption that she intends to say that Peter's bat is too big. That is, one of the assumptions which the speaker intends the hearer to recover is (55), which is an appropriately completed version of the assumption schema (53).⁵¹

(55) The speaker has said that Peter's bat is too big.

In this situation, as is also clear, it is the speaker by whom the thought interpreted by the utterance is entertained as a description of an actual state of affairs. But we have also looked at cases where the thought interpreted is entertained as a true description of some state of affairs, not by the speaker herself, but by some other person whose thought is being interpreted by the speaker. For instance, in the case of the echoic utterance presented in (47b) the speaker, Mary, has said that *P*. That is, the interpretation of this utterance also involves recovering the following description.

(56) The speaker has said that the Joneses aren't coming to the party.

But in the situation outlined for (47b) it is clearly not the speaker Mary herself by whom the thought interpreted by the utterance is entertained as a true description of some state of affairs. Rather, the utterance is her interpretive expression of a thought held by some other speaker, namely Peter, as a true description of some state of affairs.

How is the hearer to know whether the speaker who produces an utterance with a propositional form *P* is saying that *P* and not, for instance, telling him to *P* or asking whether *P*? That is, how is the hearer to know whether the recovered propositional form *P* should be or should not be integrated into the assumption schema (53)? According to Sperber and Wilson, the recovery of this assumption is not the result of a simple decoding process, carried out on the basis of the syntactic type of the sentence uttered. For one thing, they (1986:247) argue that the correlation between syntactic sentence types - such as declarative, imperative and interrogative - and generic speech acts cannot be maintained. Moreover, they suggest that even the claim that there is "a well-defined range of mutually exclusive sentence types" is open to question. They (1986:247) then continue as follows:

⁵¹ The propositional form recovered for utterance (16) is informally represented in (21), but even that informal representation is unwieldy. As ad hoc shorthand we prefer to use here, instead, simply "Peter's bat is too big". (Strictly speaking, of course, we should complete the assumption schema (53) by integrating into it the full content of the propositional form of utterance (16).)

"What undeniably exists is not a well-defined range of syntactic sentence types but a variety of overt linguistic devices - e.g. indicative, imperative or subjunctive mood, rising or falling intonation, inverted or uninverted word order, the presence or absence of Wh-words, or of markers such as 'let's' or 'please' - which can guide the interpretation process in various ways."

Consequently, the use of terms such as *declarative sentence* and *imperative sentence* should be seen as "nothing more than a convenient shorthand". It is against this background that the reader should interpret Sperber and Wilson's (1986:247-248) answer to the question of how the hearer can know that the speaker is saying that *P*:

"When you say that *P*, you communicate that you are saying that *P*. You may communicate this by means of explicit indicators such as indicative mood, declarative word order and so on; in the absence of such indicators, as in telegraphic forms of writing and speech, it is up to the hearer to decide whether the speaker is saying that *P* or performing one of the other generic speech acts. *In this, as in every other aspect of interpretation, he should adopt the first assumption that is consistent with the principle of relevance* [the italics are ours - M.S.-W.K.W.]."

For our purposes, the last (italicised) comment is particularly significant. It means that the hearer faced with the task of deciding whether the speaker intends to say that *P* or to tell him to *P* will again be guided by the task-specific criterion (11): he should select the first accessible interpretation, and test this for consistency with the principle of relevance. The speaker who intends to communicate that she is saying that *P* must therefore ensure that the cues which she provides are such that the assumption that the speaker is saying that *P* is the first assumption accessible to the hearer.

Suppose now that the hearer has recovered the assumption that the speaker is saying that *P*. How then is the completed assumption schema (53) to achieve relevance? Sperber and Wilson (1986:249) claim that there are a variety of ways in which such a completed assumption schema can be relevant:

"...some will have the effect of an ordinary assertion, others the effect of a report of speech or thought, others the effect of an irony or dissociation, others the effect of a speech-act classification and so on."

With the exception of the last one, all the other possibilities mentioned by Sperber and Wilson have been illustrated in 4.3-4.7. Let us recap very briefly.⁵²

We have already explained that, in the case of our cricket example (16), the hearer integrates the propositional form of the utterance into the assumption schema (53), thus recovering the description (55).⁵³ In the various situations outlined in 4.3, this description provides the hearer with evidence for the further assumption (57).

⁵² For an illustration of all the various possibilities with different utterances of the same sentence cf. S&W 1986:248-249.

⁵³ For this use of the phrase *recover a description* cf. e.g. S&W 1986:248.

(57) The speaker believes that Peter's bat is too big.

And, given that he trusts the speaker enough, this in turn provides the hearer with indirect evidence for (58).

(58) Peter's bat is too big.

An utterance which achieves relevance in this way is an ordinary assertion. An ordinary assertion results when a speaker produces an utterance which is a strictly literal interpretation of some thought of hers. In the case of a (strictly) literal utterance, as explained in 4.6, all the implications of the speaker's thought are also implications of the propositional form of the utterance. The hearer can then proceed to derive some further contextual implications from the assumptions represented in (57) and (58). We illustrated in 4.4 how this might work in the case of our cricket example (16). Other examples of ordinary assertions which achieve relevance in the way outlined here include the utterances (26b), (36b), (36c), (37b), (38b), (39b), (40b), (42c) and (48b), as interpreted in the various situations sketched above.

In the case of a metaphorical utterance, the speaker also intends to say that *P*, where *P* represents the propositional form of the utterance. For instance, in the case of the metaphorical utterance (43) discussed in 4.6 the speaker also intends the hearer to integrate the propositional form of the utterance into the assumption schema (53) and thus recover the description (59).

(59) The speaker is saying that the room is a pigsty.

The description, however, unlike the description for an ordinary assertion such as (16), does not provide the hearer with evidence that the speaker believes that *P* in a strictly literal sense: that the room is a pigsty. Consequently, (59) does not provide the hearer with any indirect evidence that the room is a pigsty. Rather, this description provides the hearer with some evidence for the assumption (60), as we explained in 4.6.

(60) The speaker believes that the room is untidy and dirty beyond the norm.

And, if the hearer trusts the speaker enough, (60) in its turn provides him with indirect evidence for (61).

(61) The room is untidy and dirty beyond the norm.

In this case, as explained above, the utterance is a less than strictly literal interpretation of a thought of the speaker's: only some of the implications of the speaker's thought are implications of the utterance which represents this thought. Two instances of saying that *P* which achieve relevance in a similar way to (43) are (42b) and (46).

In 4.7 we outlined Sperber and Wilson's account of the ironical utterance (49b). In this case, too, the utterance's propositional form is to be integrated into the assumption schema (53), the following description being thus recovered:

(62) The speaker has said that it is a nice day for a picnic.

This description, however, does not provide the hearer with any evidence that the speaker believes that *P* nor, consequently, with any indirect evidence that *P*. Rather, the description (62) achieves relevance by providing the hearer with evidence about the speaker's attitude to *P*:

(63) The speaker believes that it is ridiculous to say that it is a perfect day for a picnic.

If he trusts the speaker enough, (63) then provides the hearer with indirect evidence for (64):

(64) It is ridiculous to say that it is a perfect day for a picnic.

Other earlier examples of utterances which achieve relevance in the way just outlined are (50) and (51).

To illustrate the last possibility mentioned by Sperber and Wilson (1986:249) - namely, achieving relevance by means of a speech-act classification - let us consider an utterance of the sentence (65).

(65) The bus is leaving.

Suppose that, on hearing this utterance, the hearer (correctly) infers that he is to integrate the utterance's propositional form into the assumption schema (53), and suppose that by doing so he recovers the description (66).

(66) The speaker has said that the bus is leaving.

Imagine, in addition, that there has been an argument, with the speaker of (65) maintaining that the bus will not leave for another ten minutes and the hearer insisting that it will leave immediately. Imagine now that the bus does move off immediately, proving the hearer to be right and provoking the speaker to utter (65). How is the utterance relevant? Not by providing the hearer with any indirect evidence that *P*; it is already strongly manifest that *P*. Rather, (65) achieves relevance by providing the hearer with evidence for what Sperber and Wilson (1986:249) call "higher-level descriptions" such as (67) and (68).

(67) The speaker acknowledges that the bus is leaving.

(68) The speaker admits that she was wrong.

It is to such descriptions that Sperber and Wilson (1986:249) refer as "speech-act classifications".

Clearly, there are many different ways in which an assumption that the speaker has said that *P* can achieve relevance, only one of which is by having the effect of an ordinary assertion. How then is the hearer to know which particular effect the speaker intends? The answer, of course, relies on the criterion of consistency with the principle of relevance (Sperber and Wilson 1986:249):

"A speaker who wants to achieve some particular effect should give whatever explicit cues are needed to ensure that the interpretation consistent with the

principle of relevance is the one she intended to convey. Thus, when an utterance is interpreted as an ordinary assertion, this is not the result of the operation of some maxim of quality or convention of truthfulness, but simply of an interaction between the form of an utterance, the hearer's accessible assumptions and the principle of relevance."

This completes our presentation of Sperber and Wilson's sketch of their own proposal for the analysis of the speech act of *saying that P* (or, equivalently, for the analysis of declarative utterances).

We turn now to Sperber and Wilson's (1986:247) sketch of their own proposal for the analysis of the speech act of *telling the hearer to P* (or, equivalently, for the analysis of imperative utterances). Suppose the hearer realises that the speaker intends to tell him to *P*. Then the recovered propositional form of the utterance will be integrated into the following assumption schema:

(69) The speaker is telling the hearer to *P*.

Sperber and Wilson analyse *telling the hearer to P* as communicating that the thought of which *P* is an interpretive expression is entertained as a description of a desirable state of affairs. Accordingly, there is a fundamental difference between saying that *P* and telling to *P*: in the former case, as explained above, the thought interpreted by *P* is entertained as a description of an actual state of affairs. How can the hearer know that the speaker intends to tell him to *P*? That is, how can the hearer know that he should recover an assumption of the form (69)?

The essence of the answer is the same as that provided by Sperber and Wilson (1986:248) for the recovery of the assumption that the speaker intends to say that *P*, although the particulars of the mechanisms involved are of course different. In the case of telling the hearer to *P* the main linguistic device which can point the hearer in the right direction is imperative sentence structure. But here too there is no simple decoding procedure at work: the hearer must adopt the first accessible interpretation consistent with the principle of relevance, and it is the speaker's responsibility to ensure that this is indeed the intended interpretation.

Once the hearer has recovered the assumption that the thought of which *P* is an interpretive expression is entertained as a description of a desirable state of affairs, he must proceed with the interpretation of the utterance. Sperber and Wilson (1986:251) make some comments on this further process:

"Who entertains the thought in this way: the speaker or someone whose thought the speaker is interpreting? From whose point of view is the state of affairs described desirable? The hearer has to answer these questions inferentially. As usual, the first interpretation consistent with the principle of relevance will be selected, and a speaker who wants to be correctly understood must make sure that the interpretation she intends to convey is the first one consistent with the principle of relevance."

The answers to the two questions considered in this quote will be implicatures of the utterance in question. Their recovery is therefore in all respects identical to the recovery of the implicatures of ordinary assertions outlined in 4.4.

Sperber and Wilson (1986:250-251) suggest that the difference between requesting and advising reduces to different answers to the second one of the questions considered in the quote above: "From whose point of view is the state of affairs described desirable?" When *requesting*, the speaker represents the state of affairs described in *P* as desirable from her own point of view; when *advising*, she represents the state of affairs described in *P* as desirable from the hearer's point of view. From the rest of Sperber and Wilson's remarks it should be clear that choosing between a requestive and an advisory interpretation of an utterance is an inferential task, guided by the consistency criterion (10). And more generally: deciding on the intended illocutionary force of an utterance is an inferential task, guided by the consistency criterion. To conclude our discussion of the recovery of the intended illocutionary force of an utterance within a relevance-theoretic framework, let us briefly illustrate this last point with two possible (overall) interpretations of the sentence in (70), first presented as (11) in 2.4.6.

(70) Clean up your room.

Case 4.H

Imagine that (70) is uttered by a mother to her son in the following situation: the son is responsible for keeping his room in order. *The mother enters the son's room, and discovers that the room is in a terrible state. Having recovered the assumption that his mother is telling him to clean up his room, the son will resolve the above-mentioned questions as follows.* In this situation, firstly, the most accessible assumption as to whose thought the mother is representing is that it is her own thought. The situation as sketched furnishes no indication, for instance, that the mother is interpreting a thought of some third party.⁵⁴ Secondly, the most accessible assumption about who views as desirable the state of affairs referred to is that it is the mother who holds this view. The inference that it is indeed the mother who regards the state of affairs as desirable will be based on some easily accessible information about the mother's views on neatness in the house. Once the son has recovered the assumption that his mother intends to convey the information that she regards it as desirable that he should clean up his room, he can infer that she expects him to clean up his room. That is, he must interpret her utterance as a request to clean up his room. In the derivation of this implicature, the assumption that the mother regards it as desirable that her son should clean up his room will combine with another easily accessible assumption - namely, that children are expected to carry out certain actions which their parents regard as desirable. For children to clean up their rooms normally falls within this class of actions. In

⁵⁴ Let us outline a situation in which the most accessible assumption is that the speaker is interpreting someone else's thought: A teacher and some students are in a room. The teacher says something, but one of the students fails to catch her words. The student then asks one of the other students what the teacher has said. This student then responds with (70). In such a situation the most accessible interpretation would be that the speaker is interpreting a thought not of her own, but of some third person, specifically the teacher.

the situation as sketched, the requestive interpretation will therefore meet the consistency criterion.

Case 4.I

Suppose now that (70), *Clean up your room*, is uttered by Mary in response to repeated complaints by her brother, Peter, that he cannot find certain things in his room. Suppose further that Peter's room is very untidy and messy. In such a situation, (70) will be intended, and interpreted, not as a request to Peter to clean up his room, but rather as a piece of advice, or a suggestion, for him to do so. Having recovered the assumption that Mary is telling him to clean up his room, Peter will infer that it is Mary who is entertaining the thought in question, not someone else. Moreover, he will infer that she is representing the state of affairs as desirable from his, not her, point of view. This further inference will be based in part on the readily accessible assumption that things are easier to find in an ordered environment than in a chaotic environment. Note also that no further implicature that Mary is expecting Peter to clean up his room will follow in this situation. Here, then, the advisory interpretation is the one consistent with the principle of relevance.

In sum: recovering the intended illocutionary force of an utterance is always an inferential process, subject to the criterion of consistency with the principle of relevance. Illocutionary-force indicators, such as declarative and imperative mood, merely point the hearer in the direction in which he should search for the relevance of the utterance. As in all other tasks of the interpretation process, this search for relevance is guided by the consistency criterion.

4.9 Explaining utterance interpretation: the common pattern

In the course of this chapter we have considered in some detail a number of situations in which a hearer is faced with the task of correctly interpreting a linguistic utterance addressed to him by a speaker. In terms of relevance theory, a common pattern underlies the explanation of all such cases of verbal communication. It may be useful, at the end of the present chapter, to take a look at the ways in which Wilson and Sperber (1987b) characterise this common pattern.⁵⁵

In 4.3 we reviewed a number of cases illustrating how the hearer's identification of the intended explicit content (or, equivalently, the intended propositional form) of an indeterminate utterance is explained by relevance theory. The pattern common to the relevance-theoretic explanation of all such cases is characterised by Wilson and Sperber (1987b) in essentially the terms of (71), presented on page 86 below. Moreover, the same pattern underlies the resolution of indeterminacies of context and contextual effects, extensively illustrated in 4.4 to 4.8 above. Sperber and Wilson characterise this second pattern in essentially the terms of (72), which we present on page 87 below. Note that in both (71) and (72) the term *proposition* is used as a synonym for Sperber and Wilson's term *propositional form* (explained and illustrated in 4.3 above).

⁵⁵ In the characterisations presented below, we follow closely W&S's (1987b:156) own wording, though with some extensions and a different lay-out.

Commenting in general terms on their book *Relevance. Cognition and communication*, Sperber and Wilson (1987b:156) claim that, in it, they

"... show ... that in every aspect of utterance interpretation involving the resolution of indeterminacies - of content, context or contextual effects - the first accessible interpretation consistent with the principle of relevance is the only assumption consistent with the principle of relevance, and is the one a rational hearer should choose."

If the foregoing presentation of Sperber and Wilson's relevance-theoretic account of the various aspects of utterance interpretation has enabled our reader to gain a clearer idea of what they mean by this claim, then chapter 4 will have achieved its overall purpose.

(71) *Resolving indeterminacies of content*

Part A: The first accessible interpretation which yields adequate contextual effects is consistent with the principle of relevance.

Consider a situation in which a hearer is interpreting a linguistic utterance addressed to him by a speaker. Assume that this situation also has two further properties. First, the utterance may be taken to express either of two propositions, P1 and P2. Second, it is manifest in the shared cognitive environment of speaker and hearer

- (i) that P1 is more accessible to the hearer than P2;
- (ii) that the hearer has access to a context C1 in which P1 has enough contextual effects to be worth his attention (call this set of effects "E1");
- (iii) that there is no more-readily accessible context in which P1 would have enough contextual effects to be worth the hearer's attention; and
- (iv) that there is no alternative utterance (or, more generally, act of inferential communication) which would have achieved the set of effects E1 more economically.

Then the overall interpretation which results from selecting P1 and processing it in the context C1 to obtain the set of effects E1 is consistent with the principle of relevance.

Part B: The first interpretation found to be consistent with the principle of relevance is the only such interpretation.

Moreover, the overall interpretation which results from selecting P1 and processing it in the context C1 to obtain the set of effects E1 is the *only* interpretation consistent with the principle of relevance. For,

- (i) recall that all the facts presented in Part A above are by hypothesis manifest to the speaker, and
- (ii) suppose that the speaker had nonetheless intended to express proposition P2.

Then an interpretation which results from selecting P2 and processing it in some context, however relevant that interpretation, would *not* be consistent with the principle of relevance, because recovering it would put the hearer not only to the justifiable processing effort of,

- (i) first, recovering and processing P1,

but also to the unjustifiable processing effort of,

- (ii) next, recovering and processing P2, and,
- (iii) then, engaging in some further form of inference in order to decide between the two resulting interpretations, choosing the one based on P2.

This unjustifiable processing effort the speaker could have spared the hearer by simply rephrasing her utterance to eliminate P1 entirely, or to make P2 more accessible than P1.

(72) *Resolving indeterminacies of context and contextual effects*

Part A: The first accessible interpretation which yields adequate contextual effects is consistent with the principle of relevance.

Consider a situation in which a hearer is interpreting a linguistic utterance addressed to him by a speaker. Assume that this situation also has two further properties. First, the utterance may be taken to be intended to be interpreted in either of two contexts, C1 and C2. Second, it is manifest in the shared cognitive environment of speaker and hearer

- (i) that context C1 is more accessible than context C2;
- (ii) that context C1 will combine with the proposition P expressed by the utterance to yield enough contextual effects, E1, for a small enough processing effort, to be worth the hearer's attention;
- (iii) that no other utterance would have achieved these effects more economically.

Then the overall interpretation which results from processing the expressed proposition P in the context C1 and deriving its associated contextual effects E1 is consistent with the principle of relevance.

Part B: The first interpretation found to be consistent with the principle of relevance is the only such interpretation.

Moreover, the overall interpretation which results from processing (the expressed proposition) P in the context C1 and deriving its associated contextual effects E1 is the *only* interpretation consistent with the principle of relevance. For:

- (i) recall that all the facts presented in Part A above are by hypothesis manifest to the speaker, and
- (ii) suppose that the speaker had nonetheless intended to convey an overall interpretation based on context C2 and its associated contextual effects E2.

Then an interpretation which results from processing P in the context C2 and deriving its associated contextual effects E2, however relevant that interpretation, would *not* be consistent with the principle of relevance, because recovering it would put the hearer not only to the justifiable processing effort of

- (i) first, processing P in the context C1 to derive the contextual effects E1,

but also to the unjustifiable processing effort of,

- (ii) next, processing P in the context C2 and deriving the contextual effects E2, and,
- (iii) then, engaging in some further form of inference in order to decide between the two resulting interpretations, choosing the one based on C2 and E2.

This unjustifiable processing effort the speaker could have spared the hearer by simply rephrasing her utterance to eliminate C1 entirely, or to make C2 more accessible than C1.

Chapter 5

IN CLOSING

What is the merit of Sperber and Wilson's relevance theory as a new approach in the field of pragmatics? So far, in presenting relevance theory, we have focussed on its content and have not dealt explicitly with factors bearing on its merit. In closing, however, we would like to note some factors that do relate to this question, and would like to say how we stand on this question in view of those factors.

On the one hand Sperber and Wilson seem to make some sweeping claims about their relevance theory. As we have tried to illustrate in 4.3 to 4.8 above, they claim that their theory "applies to every aspect of utterance interpretation" (Wilson and Sperber 1986:77; 1987:15-16). Indeed, they claim, in more general terms, that the principle of relevance (which is the heart of their theory) is "enough on its own to yield an explanatory pragmatic theory" (Wilson and Sperber 1986:72; 1987:14; cf. also Sperber and Wilson 1986:vii, 162-163, 170, 248, 251).

On the other hand, they occasionally also seem to state their claims in far less confident terms. Recall, for instance, that in 4.8 we noted some ideas from Sperber and Wilson's (1986:243) "sketch" of their own "proposals" for the analysis of certain speech acts. In the course of their sketch they (1986:251) observe that they "believe that along these lines a satisfactory account... might be constructed". What we find striking is that the authors here expressly present their ideas as exploratory and tentative; these ideas are presumably not, in the authors' opinion, anything like a final answer. Elsewhere too the authors (Sperber and Wilson 1986:117,170; Sperber and Wilson 1987:709-710) admit frankly that some of the answers they put forward are tentative and, moreover, that at several points non-trivial questions are being left unanswered. Most modest of all, perhaps, is the opinion reflected by words such as *speculative* and *general* in the *that* clause of the following remark (Sperber and Wilson 1987:709-710), a remark addressed primarily to the readers of *Behavioral and Brain Sciences*: "We are well aware that the view developed in *Relevance...* is very speculative and, as it stands, too general to determine directly either specific experiments or computer simulations."¹

It is not surprising, then, that relevance theory has prompted the publication of a large number of responses of various kinds. Some of them have been critical discussions, both positive and negative, including the following: (Austin 1987; Kempson 1988b, 1988c, 1988d; Levinson 1989; Mey and Talbot 1988; Nemo 1988; Pateman 1986; Seuren 1987;

¹ Interestingly enough, by contrast, Kempson (1988c:vii), speaking from the viewpoint of her investigations into quantification and the form of semantic representation, considers that relevance theory, "despite the very great generality of its central claims, impose[s] a surprising number of precise consequences on the form of grammar which seem... to be correct".

Ziv 1988), and also the various contributions to the Open Peer Commentary in *Behavioral and Brain Sciences* 10 (1987): 710-736).²

Even a cursory glance at these critical discussions makes it plain that so far there is no consensus about the merit of relevance theory as a new approach in the field of pragmatics. The absence of such a consensus may be illustrated by means of the following example, involving Kempson and Levinson. Kempson (1988c:vii-viii) considers relevance theory to be characterised by "real elegance" in the sense that "precise consequences continue to pour from its ...principles", and she is thankful that these consequences have been to her "a continuing stream of intellectual stimuli demanding actions in response". Elsewhere she (1988d:160) expresses her favourable assessment of relevance theory in more general terms, including the following:

"...relevance theory ...articulates a detailed explanation of processes involved in utterance interpretation with predictions which are precise enough to bring pragmatics back into the field of serious inquiry..."

In stark contrast to Kempson's favourable assessment of relevance theory is Levinson's (1989:462) unfavourable judgement:

"...the theory to my mind never quite gets off the ground..."³

The absence of consensus illustrated by this example is perhaps familiar enough in the case of a new theory, however, and is even less surprising where the new theory is perceived as representing a new approach to an entire field.^{4,5} Debates about the merits of such theories

² Sperber and Wilson, in their turn, have replied to some of these responses. Cf. e.g. their Author's Response in Volume 10 (1987:736-751) of *Behavioral and Brain Sciences*, and also W&S 1987b.

³ More generally there is of course a lack of consensus amongst theorists of "language and mind", including not only pragmatists but also formal semanticists, formal syntacticists and philosophers of language, as is duly noted by Kempson (1988a:3-4) in her introduction to a recent collection of papers by linguists. She (1988a:3) presents this introduction as containing "a sketch of the state of the art with respect to current theories of language and mind". And then, towards the end of her introduction, she (1988a:21) observes that "the present state of the art is in general unrelentingly tribal as the chapters of this volume display, each researcher articulating and evaluating solutions within the confines of their own selected paradigm [our emphasis - M.S.-W.K.W.]".

⁴ Sperber and Wilson (1986:vii) themselves claim that they present "a new approach to the study of human communication". This view is shared by Levinson (1989:469), for instance, a pragmatist who is highly critical of relevance theory.

⁵ Bunge (1967:346) makes the point that "there is hardly a scientific field where an important theory reigns undisputed or, at least, where no alternative theories are conceivable. The more important a theory the more rivals it is likely to have". The possibility, and indeed the existence, of rival theories in most scientific fields is explained in part by the nature of the criteria by which scientists evaluate scientific theories. These criteria are numerous and diverse: alongside empirical criteria, Bunge (1967:351-356) for instance identifies twenty non-empirical criteria which he groups into five categories (formal, semantic, epistemological, methodological and metaphysical). Yet even so these criteria are *not* "truth conditions" or "decision criteria". Bunge (1967:356) goes on to emphasise that the most "that can be secured is a set of numerous and partly independent controls, which are singly insufficient to guarantee complete truth but which can jointly detect partial truth", and that all that these controls do "is to show the extent to which any factual theory succeeds and the extent to which it fails".

tend to be lengthy and complex affairs.⁶ In the case of relevance theory, moreover, the debate would be further complicated if the participants disagreed about the criteria (and/or other considerations) by means of which to assess the relative merits of pragmatic theories. As we read him, Levinson, in his review of *Relevance* (1989:469-470), points out at least one such disagreement. On the one hand, according to him, pragmatists like Sperber and Wilson place a high value on a theory in which a variety of phenomena of meaning are reduced to a single explanatory principle; on the other hand, many pragmatists (including himself) are "infinitely suspicious" of such simplicity in the theory of meaning. For pragmatists such as Levinson (1989:469), recent progress in the theory of meaning has been characterised by "the discovery of lots of different aspects of meaning, each with its different properties".

One faces here an important question: the question of criteria for theory assessment. By what criteria do, or should, pragmatists assess the relative merits of pragmatic theories? This is not the place to attempt a detailed answer to this question; yet, clearly, such criteria and/or other considerations are important as a class of factors relating to the assessment of the merit of relevance theory. So before we turn to consider a second category of published responses to relevance theory, let us pause to have a brief look at remarks by Sperber and Wilson (1987:709-710) and by Levinson (1989:456,461-465,469-470) which seem to reflect elements of partial answers to the question of criteria (and/or other considerations) by which to assess the merits of pragmatic theories.

Sperber and Wilson (1987:709-710), as noted above, concede that their 1986 version of relevance theory is "too general" to lend itself "directly" to assessment by experiment or computer simulation. They (1987:710) proceed, however, to point to - what we presume to be - criteria of assessment other than those which necessitate experimental tests and/or computer simulations. These other criteria, they (1987:710) assert, "should be kept in mind" in "assessing a new approach" to the study of "human communication". They point to these other criteria by stating a number of questions. But they do not document, illustrate, explicate or otherwise comment on the criteria of assessment which they point to by stating these questions. These other criteria, it seems to us, are more numerous and more diverse than Sperber and Wilson's questions may at first sight suggest. In a provisional attempt to bring out more clearly these two properties of these other criteria, we now partially restate Sperber and Wilson's (1987:710) questions as follows (where, from the latter half of (i) onwards, *approach* is short for *approach to the study of human communication*):

⁶ Let us note here two of Bunge's (1967:347-351) examples of "important scientific controversies" extending over long periods. Firstly, there is the Ptolemy-Copernicus dispute: the dispute about a geostatic as against a heliostatic "system of the world". This dispute arose in the earlier half of the sixteenth century but in the view of "eminent philosophers has not yet been settled" (Bunge 1967:347-348). In Bunge's (1967:348-349) own opinion, "although in the beginning the two theories were empirically equivalent, empirical differences were worked out in time and, from the start, the two theories were conceptually inequivalent and were backed by conflicting philosophies". Secondly, there is "the strife over Darwin's theory of the origin of species, which triumphed over its two main rivals, creationism and Lamarckianism, after a long and bitter struggle which in some quarters is hardly over" (Bunge 1967:349).

- (1) (i) How does the new approach to the study of human communication compare with current approaches in regard to "explicitness"?
- (ii) How does the new approach compare with current approaches in regard to "plausibility"?
- (iii) How does the new approach compare with current approaches in regard to "generality"?
- (iv) How does the new approach compare with current approaches in regard to "explanatory power"?
- (v) Does the new approach "throw new light on the very rich and diverse data available to all of us as individuals involved in communication"?
- (vi) Does the new approach "throw new light on the narrower but more reliable data gathered by scholars"?
- (vii) Does the new approach "suggest new empirical research"?
- (viii) Is the new approach "relevant to more than one of the many disciplines involved in the study of human communication - linguistics, pragmatics, philosophy, cognitive psychology, artificial intelligence, social psychology, literary studies, anthropology, and sociology"?
- (ix) Could the new approach "foster fruitful interactions" among the above-mentioned "disciplines involved in the study of human communication"?

Note in particular that Sperber and Wilson (1987:709-710) do not state any considerations relating to the mode of application of the criteria which they point to - for instance, should the criteria be applied in a random sequence or in some fixed sequence? - or to the relative importance of the criteria or, indeed, to any kind of interconnectedness among the criteria.

In his review of *Relevance*, Levinson (1989) does not expressly list, state or point to any criteria (and/or other considerations) by which, in his opinion, (pragmatic) theories are to be assessed. Nevertheless, various remarks which he makes in the course of his review do seem to be based on criteria, and other considerations, of just this sort. Some of these are pointed to in the list of statements below:

- (2) (i) The criteria for theory assessment include some criterion of "coherence".⁷
- (ii) The criteria for theory assessment include some criterion of "clarity".⁸
- (iii) The criteria for theory assessment include some criterion of "explanatory power".⁹

⁷ Cf. e.g. Levinson 1989:466, and, in particular, the mention of "coherence": "the theory to my mind never really gets off the ground for two reasons: the first is lack of internal coherence and clarity".

⁸ Cf. the comment quoted in n.7 above, and, in particular, the mention of "clarity".

⁹ Cf. e.g. Levinson 1989:466, and, in particular, the phrase "that R can explain": "The reader thus comes away from each sample analysis with the distinct impression of sleight of hand, not the impression of a demonstration that R [=the principle of relevance - M.S.-W.K.W.] can explain the inferences in question."

- (iv) Explanatory power as a property of a theory requires the theory to be "coherent" and "clear".¹⁰
- (v) The criteria for theory assessment include some criterion of "falsifiability".¹¹
- (vi) The criteria for theory assessment include some criterion of "predictive power".¹²
- (vii) The criterion of "falsifiability" is in some sense less demanding than the criterion of "predictive power".¹³
- (viii) The criterion of "falsifiability" and also the criterion of "predictive power" are in some sense less demanding than the criterion of "explanatory power".¹⁴
- (ix) Factors postulated in a theory must be "plausible" in the sense of being supported by, rather than in conflict with, existing evidence.¹⁵
- (x) The criteria for theory assessment contain some criterion by which a theory is the better, the more closely attuned it is to pre-theoretical intuition.¹⁶

It is not evident from Levinson's (1989) text whether he subscribes to any specific ideas about the mode of application of criteria such as those pointed to in, for instance, (i) and (x) immediately above, or about the relative importance of such criteria. If we are not mistaken, then (iv), (vii) and (viii) immediately above do indicate that he does subscribe to certain ideas about the interconnectedness of at least some of the criteria for theory assessment.

¹⁰ Cf. e.g. Levinson 1989:463, where "the first [problem - M.S.-W.K.W.]" refers back to an alleged lack of coherence and clarity in relevance theory: "The second main problem derives from the first. If the account is to be of any use, R [=the principle of relevance - M.S.-W.K.W.] must make clear predictions, or at the very least it should be possible to falsify R. But unfortunately R does not seem to have that kind of clear application".

¹¹ Cf. the comments quoted in n.10 above, and, in particular, the phrase "it should be possible to falsify R".

¹² Cf. again the comments quoted in n.10 above and, in particular, the phrase "R must make clear predictions".

¹³ Cf. again the comments quoted in n.10 above, and, in particular, the words "or at the very least". Presumably, by *internal clarity* Levinson means 'conceptual clarity'.

¹⁴ This idea is implicit in the middle sentence of the following comment by Levinson (1989:466), where "the first [problem - M.S.-W.K.W.]", as we point out also in n.10 above, refers back to an alleged lack of coherence and clarity in relevance theory: "The second main problem derives from the first. If the account is to be of any use, R [=the principle of relevance - M.S.-W.K.W.] must make clear predictions, or at the very least it should be possible to falsify R. But unfortunately R does not seem to have that kind of clear application."

¹⁵ Cf. e.g. Levinson 1989:456: "The book ... ignores many current developments in semantics, pragmatics and the study of inference" and (1989:464): "...what *a priori* reason is there to think that these two factors E and C [effort and cost, respectively - M.S.-W.K.W.] are the crucial factors determinative of utterance interpretation? E is a very specific measure of informational richness; most pragmatic theorists assume two such measures, incommensurable, indeed intrinsically opposed to each otherC is a measure of inference-cost; yet all the evidence points to inference being 'cheap' and profligately employed...".

¹⁶ Cf. e.g. Levinson (1989:467): "A further problem is that the pre-theoretical notion of relevance is not well captured by R [=the principle of relevance - M.S.-W.K.W.]..." and Levinson (1989:467): "There are ... reasonably well worked out alternatives to the theory of [relevance - M.S.-W.K.W.]: theories that attempt to capture some of the content of pre-theoretical judgements of relevance... Each of these is closely attuned to intuition...".

The two lists above are highly tentative; they do not result from the kind of large-scale and in-depth analysis that would be needed to produce a comprehensive and precise account of the criteria for theory assessment subscribed to by Sperber and Wilson or by Levinson. These two lists merely reflect *elements of partial* answers to the question of criteria (and other considerations) by which to assess the merits of pragmatic theories. We say "partial", because obviously the lists are not necessarily complete. Moreover the terms in which, in these lists, we point to these criteria are on the whole quite vague. From these lists, nonetheless, we gain two initial impressions regarding these criteria. First impression: seemingly, these pragmatists subscribe to a large number of diverse criteria for theory assessment. If this is in fact so, the situation is comparable with that presented by Bunge (1967:346-357) as normal in scientific inquiry.¹⁷ Second impression: seemingly, the respective sets of criteria subscribed to by Sperber and Wilson and by Levinson are not identical. For instance, Sperber and Wilson subscribe to criteria of "relevance to related disciplines" and "fostering fruitful actions among related disciplines", but Levinson seemingly does not. Seemingly, too, Levinson sets greater store by some criterion of "falsifiability" than do Sperber and Wilson. The more accurate our second impression is, the more likely it is that the debate about relevance will be complex and protracted.

All in all, then, our impression is that the debate about relevance theory as an approach to the study of human communication will continue for some years yet and will not be resolved, if at all, in any straightforward way.

Along with critical responses to relevance theory there have also been responses which have taken the form of descriptive applications of the theory and/or extensions of its content, among them the following: (Blakemore 1987, 1988; Blass 1989; Carston 1988; Clark 1989; Downes 1988; Furlong 1989; Gutt 1989, 1990, 1991; Haegeman 1989; House 1989; Kempson 1986, 1988b, 1988d, 1988e; O'Neill 1989; Pilkington 1989; Smith 1988; Smith and Smith 1988; Zegarac 1989). Clearly, scholars from diverse disciplines are being stimulated to use the concepts and insights of relevance theory to seek answers to new questions and also to seek new answers to old questions. Such investigations within a relevance-theoretic framework will lead, we expect, not only to more and better descriptive applications of relevance theory but also to further extensions and revisions of the theory's content. Two interesting instances of a revision of the theory's content are mentioned by Wilson and Sperber (1986:261, n.11).¹⁸ An interesting instance of an improved application results from the way in which Sperber and Wilson (1989:104-107) have themselves revised their analysis of verbal irony within a relevance-theoretic framework. Similarly, interesting proposals for the extension and revision of the theory's content have resulted from Gutt's (1990, 1991) exploratory application of relevance theory to translation.

¹⁷ In this connection cf. again n.5 above.

¹⁸ The first of these revisions: in an earlier version of their relevance theory, Sperber and Wilson assumed that there was a presumption of maximal relevance; in the 1986 version, however, they assume instead that there is a presumption of optimal relevance (on the latter cf. 3.5 above). The second of these revisions: in an earlier version of their theory, Sperber and Wilson assumed "that communicator and audience had to have and use knowledge of the principle of relevance"; in the 1986 version, on the contrary, they assume that communicator and audience have no such knowledge.

So, how do we stand on the question of the merit of relevance theory? We agree with Leech and Thomas (1990:203) that "It is too early for an evaluation of Sperber and Wilson's relevance theory", if "evaluation" is taken in the sense of 'overall, definitive evaluation'. In our opinion, nevertheless, relevance theory has at least one property which it is not too early to evaluate - namely, its ability to generate new questions and new potential answers. This ability of relevance theory can in our opinion be evaluated as being "already considerable".

Finally, then, our decision to write this guide must not be seen as based on a conviction that relevance theory provides "final answers". Rather, our decision has been prompted by our perception that three sorts of developments are likely to happen in connection with relevance theory as a new approach in the field of pragmatics. First, as noted above, debate about the content and the merit of relevance theory is likely to continue. Second, analytical/descriptive work within a relevance-theoretic framework is likely not only to be carried on but to increase.¹⁹ Third, this debate and this work are likely to go on generating questions and insights well worth the attention of those with a scholarly or professional interest in the description and the explanation of verbal communication.

¹⁹ Similar to our perception regarding these two lines of development is the perception expressed by Leech and Thomas (1990:204): "Relevance theory as propounded by Sperber and Wilson will no doubt be a major focus for future investigations into the nature of pragmatic meaning".

REFERENCES

- Alexander, Peter. 1969. *An introduction to logic. The criticism of arguments*. London: George Allen and Unwin Ltd.
- Austin, P. 1987. Review of Sperber & Wilson 1986. *Australian Journal of Linguistics* 7:129-137.
- Black, Max. 1952. *Critical thinking. An introduction to logic and scientific method*. (Second ed.). Englewood Cliffs, N.J.: Prentice-Hall, Inc.
- Blakemore, Diane. 1987. *Semantic constraints on relevance*. Oxford: Blackwell.
- Blakemore, Diane. 1988. The organization of discourse. In Newmeyer (ed.) 1988, Vol. IV:229-250.
- Blass, Regina. 1989. Pragmatic effects of coordination: the case of 'and' in Sissala. *UCL Working Papers in Linguistics* 1:32-51.
- Bunge, Mario. 1967. *Scientific research II. The search for truth*. Berlin: Springer-Verlag.
- Carston, R. 1988. Language and cognition. In Newmeyer (ed.) 1988, Vol. III:36-68.
- Clark, Billy. 1989. A relevance-based approach to 'pseudo-imperatives'. *UCL Working Papers in Linguistics* 1:53-73.
- Collinge, N.E. (ed.). 1990. *An encyclopedia of language*. London and New York: Routledge.
- Copi, Irving M. 1973. *Introduction to logic*. (Third ed.). London: Collier-Macmillan Ltd.
- Downes, William. 1988. Discourse and drama. In Van Peer (ed.) 1988:225-257.
- Furlong, Anne. 1989. Towards an inferential account of metonymy. *UCL Working Papers in Linguistics* 1:136-144.
- Gutt, Ernst-August. 1989. Translation and relevance. *UCL Working Papers in Linguistics* 1:75-94.
- Gutt, Ernst-August. 1990. A theoretical account of translation - without a translation theory. *Target* 2:135-164.
- Gutt, Ernst-August. 1991. *Translation and relevance. Cognition and context*. London: Blackwell.
- Haegeman, Liliane. 1989. *Be going to* and *will*: a pragmatic account. *Journal of Linguistics* 25:291-317.
- House, Jill. 1989. The relevance of intonation?. *UCL Working Papers in Linguistics* 1:3-17.
- Hyman, Larry M., and Charles, N. Li. (eds.). 1988. *Language, speech and mind. Studies in honour of Victoria A. Fromkin*. London: Routledge.

- Kempson, Ruth. M. 1986. Definite NPs and context-dependence: a unified theory of Anaphora. In Myers et al. (eds.) 1986:
- Kempson, Ruth M. 1988a. The relation between language, mind, and reality. In Kempson (ed.) 1988:3-25.
- Kempson, Ruth M. 1988b. On the grammar-cognition interface: the principle of full interpretation. In Kempson (ed.) 1988:199-224.
- Kempson, Ruth M. 1988c. Preface. In Kempson (ed.) 1988:vii-viii.
- Kempson, Ruth M. 1988d. Grammar and conversational principles. In Newmeyer (ed.) 1988, Vol II:139-163.
- Kempson, Ruth M. 1988e. Logical form: the grammar-cognition interface. *Journal of Linguistics* 24:393-341.
- Kempson, Ruth M. (ed.). 1988. **Mental representations. The interface between language and reality.** Cambridge: Cambridge University Press.
- Leech, Geoffrey, and Jenny Thomas. 1990. Language, meaning and context: pragmatics. In Collinge (ed.) 1990:173-206.
- Levinson, S. 1983. **Pragmatics.** Cambridge: Cambridge University Press.
- Levinson, Stephen C. 1989. A review of *Relevance*. *Journal of Linguistics* 25:455-472.
- Lyons, John. 1977. **Semantics. Volume 1.** Cambridge: Cambridge University Press.
- Mey, Jacob L., and Mary Talbot. 1988. Computation and the soul. *Journal of Pragmatics* 12:743-789.
- Myers, T., K. Brown, and B. McGonigle. (eds.). 1986. **Reasoning and discourse processes.** London: Academic Press.
- Nemo, Francois. 1988. Review of Sperber and Wilson 1986. *Journal of Pragmatics* 12:791-795.
- Newmeyer, Frederick (ed.). 1988. **Linguistics: the Cambridge survey. Volumes I-IV.** Cambridge: Cambridge University Press.
- O'Neill, John. 1988. Relevance and pragmatic inference. *Theoretical Linguistics* 15(3):241-261.
- Pateman, Trevor. 1986. Relevance, contextual effects and least effort. *Poetics Today* 7(4):745-754.
- Pilkington, Adrian. 1989. Poetic effects: a relevance perspective. *UCL Working Papers in Linguistics* 1:119-134.
- Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. 1976. **A grammar of contemporary English.** London: Longman.
- Seuren, Pieter A.M. 1987. The self-styling of relevance theory. *Journal of Semantics* 5:123-143.
- Smith, N.V. 1988. First language acquisition and relevance theory. *polyglot* 9/2: 2-29.

- Smith, Neil. 1989. Can pragmatics fix parameters? **UCL Working Papers in Linguistics**: 169-179.
- Smith, Neil, and Amahl Smith. 1988. A relevance-theoretic account of conditionals. In Hyman and Lee (eds.) 1988:322-352.
- Sperber, Dan, and Deirdre Wilson. 1986. **Relevance. Communication and cognition**. Oxford: Basil Blackwell.
- Sperber, Dan, and Deirdre Wilson. 1987. Précis of *Relevance. Communication and cognition*. **Behavioral and Brain Sciences** 10:697-710.
- Van Peer, Willie (ed.) 1988. **The taming of the text. Explorations in language, literature and culture**. London and New York: Routledge.
- Wilson, Deirdre, and Dan Sperber. 1986. Pragmatics and modularity. **Chicago Linguistic Society 22: Parasession on pragmatics and grammatical theory**.
- Wilson, Deirdre, and Dan Sperber. 1987a. An outline of relevance theory. **Notes on Linguistics** 39:5-24.
- Wilson, Deirdre, and Dan Sperber. 1987b. The self-appointment of Seuren as censor. A reply to Pieter Seuren. **Journal of Semantics** 5:145-162.
- Wilson, Deirdre, and Dan Sperber. 1989. On verbal irony. **UCL Working Papers in Linguistics** 1:96-117.
- Zegarac, Vladimir. 1989. Relevance theory and the meaning of the English progressive. **UCL Working Papers in Linguistics** 1:19-30.
- Ziv, Yael. 1988. On the rationality of 'relevance' and the relevance of 'rationality'. **Journal of Pragmatics** 12:535-545.

CONTENTS OF PREVIOUS ISSUES OF SPIL PLUS

SPIL PLUS 1 (1980)

Jeanne Maartens	Nuwe ontwikkelings binne Chomsky se teorie van kerngrammatika	1-78
-----------------	---	------

SPIL PLUS 2 (1980)

Alta Oosthuizen	Leksikale basiskategorieë. Probleme vir die "working grammarian"	1-217
-----------------	--	-------

SPIL PLUS 3 (1980)

Jeanne Maartens	<i>Afrikaanse Sintaksis</i> : 'n evaluering	1-91
-----------------	---	------

SPIL PLUS 4 (1981)

Walter Winckler	Aspekte van die taksonomiese taalkunde	1-166
-----------------	--	-------

SPIL PLUS 5 (1981)

Fonologie, semantiek, taalpolitiek: drie besprekings

R.P. Botha	Oor vrye kritiese meningswisseling en die skete van peripatete	1-6
J. Roux	<i>Praktiese fonetiek vir taalstudente</i> . . . : 'n Evaluering	7-43
C. le Roux en M. Sinclair	'n Kritiese beskouing van <i>Inleiding tot die Semantiek</i>	44-92
R.H. Pfeiffer	Bedenkinge oor die politieke kader van <i>Tuiste in eie Taal</i>	93-105

SPIL PLUS 6 (1982)

Repliek en kritiek

J.C. Steyn	Bedenklike Bedenkinge: Pfeiffer se Taal-politiek	1-32
D.P. Wissing	<i>Fonetiek vir Eerstejaars</i> : 'n kort waardering	33-39

SPIL PLUS 7 (1982)

Chomsky in die jare tagtig

R.P. Botha	'n Roete deur die wêreld van Chomsky	1-47
W.K. Winckler	<i>Rules and Representations</i> : Lees maar ... Daar staan wel wat daar staan	48-94

SPIL PLUS 8 (1983)

Taalwetenskap en die studie van literatuur

Rudolf P. Botha	Wat kan Taalwetenskap die student van literatuur bied?	1-24
Arnold Blumer	Basiskonsepte vir die doseer van taal en literatuur in die konteks van 'n vreemde kultuur	25-30
Wilhelm Liebenberg	Die rol van die Taalwetenskap in die ontwikkeling van die Literatuurwetenskap	31-43
Henning Snyman	Implikasie as taalverskynsel	44-51
Johan C. Thom en Pieter G.R. de Villiers	Die relevansie van 'n linguistiese benadering tot die analise van klassieke tekste	52-61

SPIL PLUS 9 (1984)

Hester Waher	Chomsky se teorie van 'government-binding'	1-187
--------------	--	-------

SPIL PLUS 10 (1985)

Afdeling I

Algemene perspektiewe op die normering van taalgebruik

D.D. Joubert	'n Sosiologiese konsep van norme	1-12
J.C. Steyn	Algemene aspekte van taalnormering	13-69
V.N. Webb	Taalnorme en Afrikaans. 'n Geval van twyfel en vertwyfeling?	70-108
W.K. Winckler	Oor die taalwetenskaplike gronde van taalgebruiksnorme: "important" of "importantly"?	108-122
R.P. Botha en M. Sinclair	Die negende gebod in taalgebruik	123-146

Afdeling II

Perspektiewe op die praktyk van die normering van Afrikaans

E. Raidt	Historiese perspektief op die normering van Afrikaans	147-168
L.C. Eksteen	Die rol van die Akademie in die standaar-disering van die Afrikaanse spelling	169-200
D.C. Hauptfleisch	Taalnormering en woordeboeke: 'n Praktyk-gerigte perspektief	201-237
T.J.R. Botha	Normering in die Afrikaanse Radio- en Televisiediens	238-268
J. Picard	Die normering van Afrikaanse vakterminologie	269-291
J.S. Schutte	Die taalpraktisyn as kruispunt	292-320
D.W. le Roux	Die Staatstaaldiens: Funksies en invloed ten opsigte van die normering van taalgebruik	321-327

G. Kroes	Die doeltreffendheid van die Afrikaanse Woordelys en Spelreëls as handleiding en naslaanbron by die onderrig van Afrikaans op tersiêre vlak	328-344
J.P. Botha	Die normering van Afrikaanse taalgebruik soos beoefen in die skool en die universiteit	345-365
J.M. Marais	Taalnormering in die skool: 'n onderwyser se perspektief	366-371
J. Esterhuysen	Die dominasie van die Afrikaner-lewensbeskouing binne die Afrikaanse spraakgemeenskap, soos gemanifesteer in taalhandboeke vir hoërskole	372-391
L.R. Heiberg	Taaldokters	392-416
A. Gardiner	Die Afrikaanse taaldokter en die Hippokratiese eed: implikasies vir die taalgebruiker	417-452
H. Snyman	Norme by literêre taalgebruik	453-463
C.J. Conradie	Die normering van Afrikaanse Bybeltaal	464-487

Afdeling III

Perspektiewe op die normering van ander tale

E. Loubser	Normering op die Nederlandse taalgebied	489-496
E. Ridge	Norms in British and South African English	497-512
S. Skorge	Norme en normkritiek ten opsigte van die Duitse taal: 'n kort oorsig	513-553
P. Brink	Taalnormering in Frankryk	554-572
J.M. Lenake & C.F. Swanepoel	Historiese aspekte van normering in Suid-Sotho	573-590
A. Wilkes	Taalgebruiksnormering in die Suid-Afrikaanse Bantoetale	591-604

SPIL PLUS 11 (1987)

Rudolf P. Botha	The generative garden game: challenging Chomsky at conceptual combat	1-283
-----------------	--	-------

SPIL PLUS 12 (1989)

Jan B. Bedaux	The portraits of Simon van der Stel, first governor of the Cape	1-18
---------------	---	------

SPIL PLUS 13 (1989)

C. le Roux	On the interface of morphology and syntax	1-417
------------	---	-------

SPIL PLUS 14 (1989)

Rudolf P. Botha	The Metaphysics Market. 1 Merchandizing language as matter	1-52
-----------------	--	------

SPIL PLUS 15 (1989)

Rudolf P. Botha	The Metaphysics Market. 2 Billing language as behavioral	1-76
-----------------	--	------

SPIL PLUS 16 (1991)

Hester Waher	Oor grense, kettings, en lisensies. 'n Sleutel tot Chomsky se teorie van grammatika	1-213
--------------	---	-------

SPIL PLUS 17 (1991)

Rudolf P. Botha	The Metaphysics Market. 3 Selling language as soul	1-105
-----------------	--	-------

THE SPIL LOGO

The logo on the front cover depicts Simon van der Stel, Dutch governor of the Cape of Good Hope from 1679 to 1699, and the founder of Stellenbosch. We have chosen to portray Van der Stel in our logo for reasons of symbolism that relate to his historical significance, his intellectual qualities, and his creole descent. Simon van der Stel was the man who, in founding the town of Stellenbosch, took a deliberate initiative towards establishing the permanency of the young Dutch settlement at the Cape of Good Hope. He has been portrayed as a man endowed with special intellectual qualities, who set great store by clear, factual thinking --- a quality which we value. His creoleness, to us, is symbolic both of the melting-pot from which emerged the South Africa of the 18th century and of the kind of future that we envisage for our country: a future unmarred by the racist divide that infects our society today. Our linking of Van der Stel's image to SPIL is intended to reflect our commitment to such a future. We are happy to note that this commitment is reflected, too, in the policy of our university, quoted below:

We reject outright all discrimination on the grounds of race, colour or creed and see ourselves as committed, unequivocally, to the dismantling of apartheid and to achieving inclusive democracy and equal opportunity for all in this fair country.