RELEVANCE AND LINGUISTIC MEANING

The semantics and pragmatics
of discourse markers

DIANE BLAKEMORE

RELEVANCE AND LINGUISTIC MEANING

The semantics and pragmatics
of discourse markers

DIANE BLAKEMORE
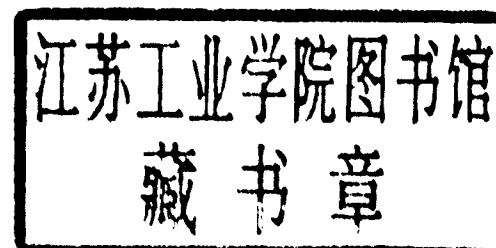
CAMBRIDGE
UNIVERSITY PRESS

**For Deirdre Wilson
with thanks**

Contents

Acknowledgements

I have never known how to thank Deirdre Wilson for the part she has played in my life – a role which goes way beyond the intellectual inspiration and support she has given me, and which has continued long after I left University College London to make my own way in the world. In dedicating this book to her, I hope to convey how important she has been to me, and to express something of the gratitude I feel.

There are many people who have helped me on my way to and through this book. Robyn Carston, with whom I have shared my half-worked-out ideas, my New Zealandness, my frustrations and moments of understanding, has shown me the importance of not thinking or working alone. The importance that her work has had for mine shows on many pages of this book. Neil Smith and Bob Borsley have continued to support and encourage me in every aspect of my work, and I would like to take this opportunity of thanking them here. Over the last year my research has gained inspiration from Corinne Iten's work on concessives: she is perhaps the only other person I know who is willing to talk about *but* at a moment's notice.

Since I have been in Manchester I could not have done without the support and friendship of Wiebke Brockhaus: she has sorted my computer problems, entertained me, listened to me and, above all, has always been there. Thanks too to Celine Berthier for making my life smoother and happier.

I would like to thank the members of the School of Languages at the University of Salford and my colleagues in the North-West Centre for Linguistics for making me so welcome in Manchester.

Last, but certainly not least, I would like to thank my family in New Zealand for their love and support, and my daughter Anna for her music, her sense of humour and her patience. I wish I could promise her that I will never talk about *but* again.

Introduction

Titles of books are often deceptive. My discovery of a copy of Austin's (1962) *How to do things with words* in a bookshop amongst various manuals on composition and writing skills was amusing but understandable. While it is unlikely that this book will be incorrectly catalogued on the basis of its title, it is possible that it will give rise to expectations that it will not fulfil. Indeed, that is its point. For the aim of this book is to show that there is no justification for writing a book about discourse or discourse markers at all.

A book which has 'discourse markers' in its title suggests that there is a class of phenomena which can be called 'discourse markers'. In earlier work (Blakemore 1987, 1992) I refer to the expressions which appear in this book as 'discourse connectives', while other writers (e.g. Fraser 1990, Schiffrin 1987, Stubbs 1983) call them 'discourse markers'. The problem is that since there is no agreement on what counts as a discourse marker, it is difficult to know whether these are two labels for the same set of phenomena. Compare, for example, the discrepancies between the lists of discourse markers given by Fraser (1990) with the one given by Schiffrin (1987):

> consequently, also, above all, again, anyway, alright, alternatively, besides, conversely, in other words, in any event, meanwhile, more precisely, nevertheless, next, otherwise, similarly, or, and, equally, finally, in that case, in the meantime, incidentally, OK, listen, look, on the one hand, that said, to conclude, to return to my point, while I have you (Fraser 1990)

> oh, well, but, and, or, so, because, now, then, I mean, y'know, see, look, listen, here, there, why, gosh, boy, this is the point, what I mean is, anyway, whatever (Schiffrin 1987)

In spite of these discrepancies, it seems that the term 'discourse' is intended to underline the fact that their role must be described at the level of discourse rather than the sentence, while the term 'marker' is intended to underline the fact that their meanings must be analysed in terms of what they *indicate* or *mark* rather than what they describe. At the same time, however, it seems to be agreed

that discourse markers are not the only expressions which operate as indicators at the level of discourse. For example, discourse adverbials such as *frankly*, *unfortunately* and *reportedly*, interjections such as *yuk* and *oh* and expletives such as *damn* and *good grief* are also described in these terms. If the term 'discourse markers' does indeed refer to a particular class of expressions, then they must have a property which distinguishes them from other discourse indicators. This property is generally considered to be their function of marking relationships or connections among units of discourse. Thus Levinson (1983) draws attention to examples of words and phrases which not only have a 'component of meaning which resists truth-conditional treatment', but also 'indicate, often in very complex ways, just how the utterance that contains them is a response to, or a continuation of, some portion of the prior discourse' (1983:87–8).[1]

It is these two assumptions – the assumption that linguistic meaning can be non-truth conditional and the assumption that there are expressions that mark connections in discourse – which this book is really about. In this way, it is less a book about expressions classified as discourse connectives or markers as a book about the theoretical assumptions that are made by those writers who have analysed expressions as discourse connectives or markers. Accordingly, it is not a book crammed with analyses of particular discourse connectives. The analyses I give (largely in chapter 4) are given in support of an approach to the study of meaning and communication in which non-truth conditionality and discourse play no role at all.

In this book I do not make a distinction between discourse markers and discourse connectives. My arguments apply to authors who use either term inasmuch as their analyses are based on the two distinctions I have just mentioned, and my own use of the terms 'discourse marker' and 'discourse connective' is not intended to reflect a commitment to a class of discourse markers/connectives. The examples that I do discuss in detail are English expressions. In this sense, I may be considered guilty of perpetuating the over-dependence of discourse marker research on English (see Schourup 1999). However, my arguments are directed at the assumptions made in discourse marker/connective research, and hence apply to any classification or analysis made on the basis of these assumptions. Nevertheless, as we shall see, my approach does raise questions about how inferential procedures are encoded in particular languages, and, moreover, sheds light on the questions about the inter-translatability of discourse markers noted by Schourup (1999). The claim that expressions such as *but* and

---

[1] For an excellent overview of the literature on discourse connectives and discourse markers, see Schourup (1999).

so are non-truth conditional has been taken to mean that their meaning must be analysed within pragmatics rather than semantics. In chapter 1, I critically examine the assumption underlying this view, namely, that semantics = truth conditions, while pragmatics = meaning minus truth conditions (see Gazdar 1979). I shall argue that this approach to the semantics–pragmatics distinction is based on a view of linguistic semantics that cannot be maintained in a cognitive approach to meaning. Semantic representations delivered by the grammar do not encode truth conditions.

Within the framework of speech act theory, the distinction between truth conditional ('semantic') meaning and non-truth conditional ('pragmatic') meaning has been unpacked as a distinction between describing and indicating. Thus it is claimed that expressions such as *but* and *so* do not contribute to the descriptive content of the utterances that contain them but merely indicate how these utterances are to be interpreted. However, as Rieber (1997) has said, while many theorists have appealed to the speech act theoretic distinction between describing and indicating, there have been relatively few attempts to say what it means for an expression to indicate information rather than describe it. The exception to this trend is Grice (1989), whose notion of *conventional implicature* has played a prominent role in discussions of non-truth conditional meaning. In chapter 2, I examine the saying–indicating distinction in detail and show how Grice's (1989) notion of conventional implicature can be regarded as following in the speech act theoretic tradition. I shall argue that, in the end, saying that an expression carries a conventional implicature simply amounts to saying it is non-truth conditional, and hence does not provide an account of what it is that such an expression contributes to.

In chapter 3, I outline an alternative approach to linguistic meaning based on Sperber and Wilson's (1995) relevance theory.[2] For Sperber and Wilson, the distinction between semantics and pragmatics is a distinction between the two kinds of cognitive processes involved in utterance interpretation. Semantic meaning is the result of linguistic decoding processes which provide an input to inferential processes constrained by a single cognitive principle. As we shall see, this approach allows for two ways in which linguistic encoding may act as input to pragmatic inferencing and hence two kinds of linguistically encoded meaning: on the one hand, a linguistic expression or structure may encode a constituent of the conceptual representations that enter into pragmatic

---

[2] References in this book will be to the second (1995) edition of Sperber and Wilson's book rather than the first (1986) edition. However, note that, apart from the postscript, the pagination of the 1995 edition is identical with that of the 1986 edition.

inferences, while on the other, a linguistic expression may encode a constraint on pragmatic inferences. This is the distinction I arrived at in my earlier book (Blakemore 1987) and which has become known as the distinction between conceptual and procedural encoding.

However, the distinction I arrived at in 1987 is not really the distinction that is at the heart of this book. I originally envisaged this distinction as being a cognitive version of the truth conditional versus non-truth conditional distinction, and hence as being co-extensive with it. Subsequent research has shown that this cannot be the case: there are expressions which encode procedures but which contribute to what is thought of as truth conditional content, and there are expressions which encode concepts but which do not contribute to what is thought of as truth conditional content. The fact that the two distinctions cross-cut each other in this way leaves us with the question of which distinction is *the* distinction in a cognitive theory of semantics. Chapters 3 and 4 of this book are intended as an argument for abandoning the distinction between truth conditional and non-truth conditional meaning in favour of the distinction between procedural and conceptual encoding.

However, in spite of the amount of research that has been inspired by the notion of procedural meaning, the notion remains relatively poorly understood. Chapter 4 addresses some of the new questions that are raised by this new approach to linguistic meaning. In particular, it addresses the question of what exactly procedural meaning looks like. As it is defined in my 1987 book, it is information about the intended cognitive effect, or, in other words, constraints on the *results* of the pragmatic inferences involved in the recovery of implicit content. However, it is not clear that this very limited notion of procedural encoding can accommodate the full range of expressions that encode constraints on relevance. The analyses in chapter 4 are not just intended as a contribution to our understanding of the roles of the expressions discussed, but aim to show how my original notion of procedural encoding must be broadened to include constraints on all aspects of inferential processing.

The picture that emerges from chapters 3 and 4 is more complex than the picture that is drawn within the speech act theoretic framework described in chapter 2. Not all the expressions which have been classified as non-truth conditional indicators indicate in the same way. Some of these expressions encode concepts, while others encode procedures. And if my analyses in chapter 3 are right, there is not just one kind of procedural encoding. All this means that from a cognitive point of view, there is not a single class of discourse markers.

This conclusion will be unwelcome not only to those who are interested in discourse markers as examples of non-truth conditional meaning, but also

to those who analyse them in terms of their function in discourse. The assumption that discourse markers operate at a discourse rather than a sentence level is based on the assumption that there is such a thing as discourse. As we shall see in chapter 5, there is more than one view of what discourse is, and accordingly more than one view of what it means for an expression to operate at discourse level. Some writers see discourse in terms of social behaviour or interaction, while others see it as an object with structural properties. On either view, discourse is an externalized object which can be investigated independently of the human mind. It is either behaviour or an abstract object. In this respect, both views must be contrasted with the one I shall argue for in this book. On this view, the object of study is not discourse, but the cognitive processes underlying successful linguistic communication, and the expressions which have been labelled as discourse markers must be analysed in terms of their input to those processes. The problem is that not all expressions classified as discourse markers make the same kind of contribution.

For those writers who do analyse these expressions in terms of their role in discourse, their most important property is their function of marking relations between discourse segments, or, in other words, their function as 'discourse glue' (Fraser 1990:385). What these expressions are assumed to connect varies according to the view of discourse that is adopted. Thus in a structural approach to discourse these expressions are analysed as marking relations between spans of discourse, while in a functional approach they are analysed as marking relations between acts or exchanges. Schiffrin (1987), who adopts a view of discourse which involves the integration of structural, semantic, pragmatic and social factors, argues that discourse markers operate on a number of different 'planes' of discourse, or that they must be analysed in terms of their role in integrating 'knowing, meaning, saying and doing' (1987:29). Building on the arguments in chapters 3 and 4, I shall argue that this view cannot be maintained, and that we can have a better understanding of the expressions which have been labelled as discourse markers if we abandon the idea that they mark connections in discourse, whether these be connections between discourse segments, connections between the propositions expressed by discourse segments or connections between social acts, and explore the idea that they contribute to relevance. This is not an argument for simply replacing the notion of discourse coherence by relevance so that we can speak of the encoding of relevance relations rather than discourse relations. Discourse, whether it is construed in structural or interactional terms, is an artifact with no psychological reality, and coherence is a property of that artifact. Relevance is not a property of discourse,

but rather of an interpretation which is mentally represented and derived through cognitive processes.

It is clear that the arguments in all parts of this book depend on the acceptance of the view that utterance interpretation involves cognitive processes: thus the overall aim of the book is not so much that there is no justification at all for writing a book about discourse or discourse connectives, but rather that there is no justification from the point of view of a cognitively grounded theory of utterance interpretation. Accordingly, it is important to explain at the outset what it means to take this approach.

The cognitively grounded theory of utterance interpretation which underlies the arguments of this book is Sperber and Wilson's (1986, 1995) relevance theory, which I outline in chapter 3. As Carston (2000a) has observed, relevance theoretic pragmatics has certain fundamental aims and assumptions in common with generative grammar. In particular, like generative grammar, relevance theoretic pragmatics aims to give sub-personal explanations rather than explanations at a personal level. This means that both theories aim at full explicitness, 'leaving nothing to the intuitions of the reader or user, so that the description or mechanisms specified could be employed by a mindless automaton with the same results as in the human case' (Carston 2000a:91). Hence the talk of computations.

It has been argued that a theory which treats utterance interpretation in this way is 'disconnected from everyday communication and its problems' (Mey 1993:82), and that the mindless automaton is an inappropriate analogy when one is trying to explain what people do when they communicate. People are 'social beings' who interact in 'pre-existing [socially determined] conditions' (Mey 1993:82). Mindless automatons are not.

Mey is right that this approach to pragmatics does not attempt to explain how people communicate. As Chomsky points out, only a theory at the personal level could explain how people communicate – how they 'pronounce words, refer to cats, speak their thoughts, understand what others say' (Chomsky 1992:213). This is not to suggest, however, that people do not operate in socially determined conditions or that people's beliefs and assumptions do not include culturally determined assumptions or assumptions about social relationships and institutions. The point is that in communicating in a social context people are enabled by various sub-personal systems – grammatical competence, an inferencing system, the visual system – and these are more amenable to scientific enquiry than the person-level activity which Mey seems to have in mind. In order to show that this view of utterance interpretation is not justified,

it would have to be shown either that communication in socially determined conditions is not enabled by a sub-personal inferencing system or that there is a theory of communication at a person-level. It seems that Mey's (1993), Leech's (1983) and Schiffrin's (1994) attempts to bridge what they call 'formalist' (or 'structuralist') approaches to language with functionalist approaches can be construed as an attempt to develop a person-level theory of communication. Thus Leech (1983) claims that 'we cannot understand the nature of language without studying both domains [grammar and pragmatics] and the interaction between them' (1983:4). In this book, I shall argue that this sort of enterprise cannot succeed since it does not explain how grammar, which according to Chomsky, is a mentally represented system, does interact with pragmatics, which according to Leech, Mey and Schiffrin, is a theory of something external to the human mind.[3]

The fact that relevance theoretic pragmatics is located in cognitive science and is aiming at sub-personal explanations means that it has the potential to provide a theory of utterance interpretation which is consistent with generative grammar. However, this should not be taken to mean that relevance theory aims to explain utterance interpretation within a grammatical model. As we shall see, according to relevance theory, communication involves two distinct cognitive mechanisms which are as different from each other as 'walking is from plane flight' (Sperber and Wilson 1995:3). If this is right, then it would be simply inappropriate to reduce communication to a single model. In particular, it would be inappropriate to explain communication within the model of generative grammar. For generative grammar can only provide an account of one of the cognitive mechanisms involved in utterance interpretation, namely, the coding–decoding mechanism. The contribution of relevance theory is to account for the other type of mechanism, which is, as I have already said, inferential.

In this respect, relevance theory differs not only from Leech's (1983) and Mey's (1993) social approaches to pragmatics, but also from the structural approaches to discourse mentioned above (for example, Mann and Thompson's (1987, 1988) rhetorical structure theory). For these approaches aim to develop a theory of interpretation within a grammatical model. Indeed, at the 1999 Conference on Economy in Language Design, Computation and Use (Lyon,

---

[3] In fact, it seems that many of these authors' examples of social interaction can be explained only if we assume that humans are able to make inferences to conclusions that are mental representations of another human's thoughts. In other words, the explanation of the fact that people can communicate in a socially determined context and affect social relationships lies in a sub-personal theory.

October 1999), pragmatic explanations were frequently praised on the grounds that they were like or analogous to grammatical explanations.

If the relevance theoretic view is right, and pragmatic explanations are *not* like grammatical explanations, then the exact nature of the relationship between the two systems has to be spelt out. Clearly, this cannot be done until it is shown how relevance theory works in actual cases of utterance interpretation. However, there are two general points which must be made in order to locate the chapters that follow within familiar theoretical paradigms in linguistics.

The first point concerns the distinction between competence and performance. According to some theorists, for example, Leech (1983) and Mey (1993), the distinction between Chomskyan linguistics and pragmatics is co-extensive with the distinction between competence and performance. Accordingly, they see the move towards an integrated theory of language (see above) as a move towards a theory in which a competence theory is complemented by a performance theory. By 'competence', both writers mean a language user's 'knowledge of the language and its rules (as e.g. described in transformational generative grammar' (Mey 1993:36)). By 'performance', they mean 'the way the individual user [goes] about using his or her language in everyday life' (Mey 1993:36). While this definition of competence in terms of knowledge is mentalistic, the definition of performance suggests socially determined behaviour.

It is clear that relevance theoretic pragmatics cannot be a performance theory in this sense. Does this mean that it belongs to the realm of competence? In chapter 1, I shall argue that there is no sense in which pragmatics can belong in the domain of semantic competence, since this is not enough to deliver the intended interpretation of an utterance. Chomsky (1980) has taken this to mean that alongside a theory of grammatical competence we need a theory of pragmatic competence or, in other words, a theory of speakers' knowledge of the conditions for appropriate use, of how to use grammatical and conceptual resources to achieve certain ends or purposes. While some writers, notably, Kasher (1991a, 1991b) have developed a view of pragmatic competence located outside the grammar, others, notably, Kuno (1990, 1987) and Prince (1985, 1988, 1997) have developed the idea that the grammar should include a pragmatic component, or in other words that there should be a theory of *linguistic* pragmatic (or discourse) competence.

Following Carston (2002), I shall argue that the idea of pragmatic competence, which has been developed by Kasher (1991a, 1991b), cannot be maintained, and that what is needed is not a competence theory, but a performance theory. This is not a theory of performance in the sense defined by Leech (1983)

and Mey (1993) since it is concerned with cognitive performance mechanisms rather than social interaction. However, nor is it a theory of linguistic performance in the sense defined, for example, by Frazier (1987), since it is not concerned with linguistic performance mechanisms. As I have underlined throughout this introduction, pragmatic processing, according to relevance theory, is inferential.

While Prince (1988) recognizes that there are aspects of utterance comprehension which fall outside linguistics, she draws attention to a range of cases in which a particular linguistic form seems to encode information about the context in which the sentence that contains it should be used. For example, the cleft construction in (1a) has the same propositional content as the cleft in (1b) and yet only (1a) is appropriate as an answer to the question in (2):

(1)   (a) It was Anna who found the money.

       (b) It was the money that Anna found.

(2)   Who found the money?

As Carston (2000a) observes, there is a certain amount of controversy about this work (see Sperber and Wilson 1995:202–17). While these particular phenomena are not the ones I am concerned with in this book, it seems that the issue that Prince has raised is of central importance to the analysis of many of the phenomena that have been treated as discourse markers. For as we shall see, these expressions also seem to encode information about the contexts in which the utterances that contain them are appropriate rather than information about their propositional content. Thus all of the utterances in (3) have the same propositional content:

(3)   (a) Anna is here. So Tom's got a meeting.

       (b) Anna is here. But Tom's got a meeting.

       (c) Anna is here. After all Tom's got a meeting.

This would seem to suggest that if Prince's phenomena are indeed to be explained in terms of a pragmatics sub-component of the grammar, so must the expressions in (3).

In this book, I argue that the explanation for these phenomena must indeed lie in the grammar, in the sense that the information that is conveyed is linguistically encoded. At the same time, I shall argue that these expressions encode information about pragmatic processes. However, these arguments assume an approach to semantics and pragmatics in which the analysis of these phenomena, although it makes reference to pragmatic information, is located in a theory of *semantic* competence.

As we shall see, this approach to the semantics–pragmatics distinction is grounded in a modular theory of mind. This brings us to the final point in the theoretical map I am attempting to draw. Generative grammar is based on a modular theory of mind in the sense that, first, the grammar is regarded as an autonomous system that does not have access to assumptions about the world, and, second, the grammar itself consists of autonomous but interacting sub-modules. Some writers have argued that pragmatics is modular in both senses. Thus Horn (1988) argues that pragmatics

> may be viewed as internally modular and interactionist, in the sense that the conceptually distinct sub-components...of pragmatic analysis may be simultaneously called upon within a single explanatory account of a single phenomenon, just as autonomous but interacting grammatical systems may interact to yield the simplest, most general, and most comprehensive treatment of some linguistic phenomenon (cf. the deconstruction of the passive in Chomsky 1982). (Horn 1988: 115)

Horn's analogy is derived from a view of pragmatics in which utterance interpretation is constrained by two distinct principles. In contrast, as we shall see, relevance theory argues that all aspects of inferential pragmatic processing are constrained by a single principle. Clearly, a theory which has only one principle cannot be modular in the second sense described above.[4]

While relevance theory is not modular in Horn's sense, it does assume a modular theory of mind in the sense that grammatical processes are different in kind from other cognitive processes, in particular, inferential pragmatic processes, and that they do not have access to the propositional world knowledge which is involved in inferential pragmatic processing. This raises the question of whether pragmatics is itself a module in the sense defined by Fodor (1983). Fodor himself argued that the human inference system is not modular but a central system which integrates and performs inferences on information derived from modular input systems (such as the visual system and the grammar) and memory. However, recently, Sperber (1994) has argued for a more radical view of modularity in which not only is pragmatics a distinct module dedicated to utterance comprehension, but also individual concepts are modules with their own inferential procedures and data bases of encyclopedic information. This argument lies outside the scope of this book (for further discussion, see Carston 2000a, 2002). My aim here has been simply to outline the theoretical assumptions underlying the discussion that will follow, and to compare them with the

---

[4] Carston (2000a, 2002) has argued that Horn's two pragmatic principles do not in fact interact in the way that grammatical principles do.

assumptions made by functional person-level approaches to language on the one hand, and cognitive sub-personal approaches, on the other.

Research on discourse connectives or markers grows daily. However, as far as I can tell, there are two distinct groups of researchers – those whose interest derives from an interest in the philosophy of language, and those whose interest derives from an interest in discourse. The two groups go to different conferences, read different literature and have different heros. Indeed, as we shall see, they have even adopted different expressions as their favourites. I am hoping that this book will be read by both groups. More particularly, I am hoping that those whose interest lies in the non-truth conditional properties of discourse connectives do not stop reading before chapter 5, and those whose interest lies in the role that these expressions play in discourse do not skip the first two chapters. The investigation of the semantics of these expressions cannot ignore their role in communication. Equally, the study of the function of these expressions in communication cannot ignore questions about the kind of information that linguistic expressions can encode.

I refer to the speaker as 'he' and the hearer as 'she'. This does not have any intended contextual implications. Examples are numbered consecutively within each chapter.

1 *Meaning and truth*

1.1 Introduction

The expressions which occupy centre stage throughout this book have played a number of different roles. In this part of the book, we shall examine the role they have played in the move towards a non-unitary theory of meaning. As we shall see, this move is itself not always a move towards the same sort of distinction, and the purpose of this and the following two chapters is to tease these distinctions apart, and to argue for a distinction between two kinds of meaning that is grounded in human cognition.

For many writers, this distinction is the distinction between semantics and pragmatics, and the significance of the expressions which I am calling discourse connectives lies in the role they have played in arguments for the existence of pragmatic meaning. Chapter 2 will examine the attempts that have been made to develop the notion of pragmatic meaning within the framework of speech act theory. This chapter focusses on the view of semantics which underlies the argument that expressions such as *but* and *well* have pragmatic meaning rather than semantic meaning.

This view is implicit in Gazdar's (1979) definition of pragmatics:

PRAGMATICS = MEANING MINUS TRUTH CONDITIONS
(Gazdar 1979:2)

According to this view, discourse connectives such as *but* must have pragmatic meaning rather than semantic meaning because they do not contribute to the truth conditional content of the utterances that contain them. And indeed, this is usually believed to be the case. For example, Rieber's (1997) analysis of *but* is based on the assumption that the suggestion that there is a contrast between the two segments of (1) is due to the presence of *but*, but that the truth of (1) depends only on the truth of the proposition in (2):

(1)   Sheila is rich but she is unhappy.
(2)   Sheila is rich and she is unhappy.

But may be a notorious example of non-truth conditional meaning, but it is not the only one. All the expressions discussed in this book will be recognized as examples of expressions which convey suggestions that are not part of the truth conditional content of the utterances that contain them. Moreover, as we shall see in the following chapter, there is a whole variety of expressions and constructions which fall outside the scope of this book which are also considered to be non-truth conditional. According to the view under investigation, the analysis of all of these phenomena cannot be provided by a semantic theory because semantics is restricted to the study of truth conditions.

Gazdar's decision to define semantics as the study of truth conditions has its origins in formal logic rather than linguistics. However, it is a decision which has been embraced by many linguists, including linguists working in the Chomskyan tradition (for example, Higginbotham 1988). In the next section we will explore the implications of bringing these two traditions together and ask whether the semantic component of a grammar, conceived in the Chomskyan sense, can indeed be truth conditional, or in other words, whether the semantic representations generated by the grammar have truth values.

Gazdar (1979) not only argues that the grammar of natural language has a truth conditional semantics component but also seems to assume that language 'has' a pragmatics (Gazdar 1979:2), or in other words, that linguistic competence includes pragmatic competence. As we have seen, this pragmatics component would include within its domain linguistically encoded non-truth conditional phenomena such as *but*. However, as Grice (1989) has shown us, the interpretation of an utterance includes information which is not part of its truth conditional content and which cannot be obtained through decoding linguistic form. For example, according to Grice (1989) the information in (3b) is an implicature derived from B's utterance in (3a) on the basis of contextual information and the assumption that the speaker is conforming to a general principle or maxim of conversation (in this case, the maxim of relation).

(3)   (a) A: Is Anna here?
             B: She's got a meeting.
         (b) Anna is not here.

Since the truth of B's utterance would not be affected by the falsity of (3b) (it would be possible to say, 'She's got a meeting but she's here'), this implicature would, by Gazdar's criterion, be included within the domain of pragmatics. This would in itself be regarded as uncontroversial. The problem is that the inclusion of this sort of phenomenon in pragmatics along with the suggestion conveyed by *but* would seem to suggest a pragmatics that is both a component

and not a component of the grammar, or, in other words, a pragmatics which spreads across the linguistic/non-linguistic boundary.

In fact, as we shall see in section 3 of this chapter, a case has been made both for a theory of pragmatic competence whose domain includes the role of the context and general pragmatic principles in the interpretation of utterances and for a theory of linguistic pragmatic competence whose domain includes the role that certain (non-truth conditional) expressions play in the interpretation of the utterances that contain them. Thus while Kasher (1991a, 1991b) has argued for a notion of pragmatic competence which is analogous to the notion of grammatical competence but whose domain is the area of interpretation that involves inference mechanisms constrained by general principles of communication, Prince (1988) has focussed on a range of linguistically encoded structures whose use seems to depend on the speaker's knowledge of the contexts in which they are appropriate. I shall argue, first, that while Kasher's principles may be pragmatic, they are not part of a theory of pragmatic *competence*, and second, that while Prince is right to say that linguistic structures may encode information about the way an utterance is interpreted in context, these expressions should not necessarily be excluded from the domain of semantic competence.

Clearly, non-truth conditional expressions can be included within the domain of linguistic semantics only if linguistic semantics is not itself truth conditional. In section 4 I discuss the implications of this position, and, in particular, Levinson's (2000) charge that it constitutes a position of 'semantic retreat'. I shall argue that Levinson's own position is itself the result of a confusion between two different conceptions of semantics, and that in divorcing linguistic semantics from truth conditional semantics one is not necessarily abandoning the idea that there are relations between representations and the real world. There *are* relations between mental representations and the world, but these are not captured within a linguistic theory or grammar.

1.2 Meaning, truth and grammar

It is generally agreed by linguists working within the generative tradition that meaning is one of the concerns of a generative grammar. Thus according to a recent textbook on generative syntax, 'grammar is not just concerned with the principles which determine the formation of words, phrases and sentences, but also with the principles which tell us how to *interpret* (= assign meaning to) words, phrases and sentences' (Radford 1997:1). However, it seems that the notion of meaning that Radford has in mind is limited to 'structural' aspects of

meaning such as the assignment of case features, and a student-reader might be forgiven for wondering whether the theory of meaning that is meant to be the domain of generative grammar has anything to do with the theory, that she has just been introduced to in her introductory semantics course, which pairs sentences such as (4) to truth conditions in so-called T sentences such as the one in (5).

(4)   Snow is white.

(5)   'Snow is white' if and only if snow is white.

Our student will have been told that a semantics which provided a T sentence for every sentence of the language would capture the intuition that language is used to talk about the world – that there is a relationship between language and the world. But what does this mean?

For philosophers such as Davidson (1984) or Lewis (1972), it means that language itself says something about the world. This means that the meaning of a sentence is captured by a T sentence of the sort in (4) while the meanings of words are analysed in terms of their contribution to the truth conditions of the sentences containing them. For these philosophers, a language is an abstract system defined independently of the minds of the people who use it, and the goal of semantics is to construct a model of the conditions in the world that would make each sentence in that system true. In other words, their theories are the product of what Chomsky (1986) calls an externalized conception of language (*E-language*).

In contrast, within the framework of generative syntax, a language is regarded as a cognitive system that is internalized in the human brain/mind – that is, as *I-language* (Chomsky 1986). The question is whether a theory which pairs sentences with truth conditions could be part of a theory of such an internalized cognitive system.

To say that the semantic component of a generative grammar is truth conditional would be to say that the semantic representations it generates are representations of states of affairs that make them true. That is, it would be to say that semantic representations are encodings of propositions. However, it is well known that the grammatically determined semantic representations are not fully propositional because they contain expressions whose reference cannot be determined independently of the context in which they are uttered – for example, expressions such as *here*, *tomorrow*, *you* and *this*.

Within the mind-external framework of formal semantics, context dependence is accommodated by extending the formal apparatus devised for assigning truth conditions to sentences so that it assigns truth conditions to

sentence-context pairs. Thus Lewis (1972) proposed a set of contextual co-ordinates (speaker, addressee, place, time, etc.) as well as a set of possible worlds so that meanings could be defined as functions from a context-possible world pair to truth values. The context, according to this approach, is simply a specification of the identity of the speaker, the audience, time and place of utterance, the identification of any indicated objcts and anything else that is needed. What is needed, in this approach, is determined by the grammatical properties of the utterance. Thus the context for the interpretation of (6) consists of a speaker index, a place index and a time index, and the result of assigning values to these indices might result in an interpretation in which (6) is true iff Diane Blakemore arrived in Edinburgh on 22 March 2000.

(6)   I arrived here yesterday.

As Lewis himself recognized, the introduction of the context into the account raised the question of how the context is chosen for the interpretation of a particular context-dependent expression:

> Consider the sentence 'The door is open'. This does not mean that the one and only door that now exists is open; nor does it mean that the one and only door near the place of utterance is open. (1972:214)

The value of *the door*, claims Lewis, must be chosen from the objects that are somehow salient on a given occasion. And accordingly, he proposes to introduce a new contextual co-ordinate – a prominent objects co-ordinate.

However, judgements about the salience or prominence of objects are highly subjective, which suggests that the interpretation of a referring expression must depend on mental factors such as the expectations regarding the things that the speaker is likely to bring to the attention of the hearer. Morever, contextual prominence is not a sufficient condition for ensuring the correct choice of context for the assignment of reference. For example in (7) (from Blakemore 1987) the fact that both the entities referred to by *Periah's recording of the Moonlight Sonata* and *the Moonlight Sonata* are made accessible and hence salient to the hearer does not ensure that the correct reference is assigned to each of the two identical pronouns:

(7)   A: Have you heard Periah's recording of the Moonlight Sonata?
        B: Yes, *it* made me realize that I would never be able to play *it*.

In a later paper (Lewis 1979), Lewis acknowledges that the salience of objects and hence the value of a referring expression cannot be determined in advance but are adjusted in the course of interpretation so that the interpretation accommodates the assumption that the utterance of the sentence is acceptable.

In other words, the hearer uses that contextual information which yields an acceptable interpretation. However, this raises the question of what makes an utterance acceptable.

This question would seem to take us a long way from the concerns of formal semantics which, according to Lewis (1972), do not include the use of language by any person or population. It also seems to take us a long way from the concerns of generative grammarians as described by, for example, Radford (1997). For generative grammar is founded on the assumption that the grammar is an autonomous cognitive system, and that, in particular, the principles of grammar are autonomous from whatever principles constrain the use of language for communicative purposes. Indeed, according to this modular view of grammar, grammatical knowledge is qualitatively distinct and sealed off from the contextual knowledge which plays a role in, for example, reference assignment. Thus while I may be said to assign the correct referent to each of the instances of *he* in (8) on the basis of my belief that Prime Ministers are often called upon to open buildings, and do not usually work in libraries at universities, I would not be said to believe the grammatical principles which provide the basis for my judgement that (9) is ungrammatical.

(8)   My brother's going to meet the Prime Minister tomorrow. He's going to open the new library at the university where he works.

(9)   * Anna believes Ben to admire herself.

Moreover, grammatical principles cannot be said to be true or false in the way that a propositional representation is. As Chomsky says, 'the question of truth, conformity to an external reality, does not enter in the way that it does in connection with our knowledge of the properties of objects' (1980:27).

This would seem to suggest that the relations of truth and reference, which for logicians and philosophers are *the* relations in semantics, could have nothing to do with linguistic semantics as it is defined in generative grammar. Nevertheless, it has been argued that truth conditions do play a role in a linguistic theory of semantic representation. For example, Higginbotham (1988) argues for a level of semantic representation based on Davidson's (1967) version of truth conditional semantics. At this level, the speaker's knowledge of the meaning of the sentence in (10a) is represented in (10b), where the lexical entry for *walk* includes the information that the verb expresses a relation *walk (x,e)* which applies to a thing and an event if the event is an event of walking by that thing:

(10)   (a) John walks slowly.
        (b) Ee (walk, (j,e) & slow (e))

This representation, argues Higginbotham, captures the fact that our semantic competence includes the knowledge that (10a) entails both (11a) and (b).

(11) (a) John walks.
      (b) Something slow takes place.

However, in contrast with Davidson, Higginbotham argues that truth conditions are not 'things that sentences of a language *have* or possible states of affairs they answer to' (1988:31). Rather they are things that speakers of a language come to know in virtue of knowing the semantic principles of their language. This suggests that semantic competence, which is mind-internal, determines truth conditions, which are mind-external. And indeed, it seems that according to Higginbotham, the things that do have truth conditions are *utterances* of sentences, which are, of course, outside the head. In other words, it seems that Higginbotham's approach is compatible with the *I-language* approach to language advocated by Chomsky, since he is not saying that semantic representations are representations of the external world, but rather that they are conditions on the truth conditions of utterances. Such a semantics, claims Higginbotham, can 'exploit the advantages of the usual truth-theoretic paradigm without running afoul of contextual entanglements' (1988: 29–30).

As we have seen, the fact that natural language sentences contain expressions whose reference depends on the context means that they cannot be said to directly encode anything that has truth conditions. Higginbotham's (1988) suggestion for getting us out of this 'contextual entanglement' is that we use a system of conditional normal forms (see Burge 1974) in order to capture speaker/hearers' knowledge of the meanings of sentences containing demonstratives and indexicals. While the antecedent of such a conditional contains the condition on truth conditions imposed by the linguistic properties of the context-dependent expression, the consequence contains the statement of truth conditions. Thus what the hearer knows when she knows the meaning of a sentence such as (12a) is represented in (12b).

(12) (a) *She* is lazy.
      (b) If x is referred to by *she* in the course of an utterance of (12a) and
          x is female, then that utterance is true just in case *lazy (x)*.
                                                              (Higginbotham 1988: 35)

Since the truth value of the antecedent of (12b) depends on the context in which (12a) is interpreted, (12b) must be construed as a specification of particular conditions that must be satisfied by the context in order for the hearer to give a specification of the truth conditions for (12a). Thus contextual entanglements

are avoided in the sense that (12b) gives no indication of whether or how these conditions are satisfied in a particular case.

If (12b) is part of linguistic competence, then, assuming the distinction between grammatical and world knowledge described above, we would not wish otherwise. However, then (12b) must be regarded as a specification of the contextual parameters whose values have to be set for the identification of the truth conditions of any utterance containing *she*, and one would expect at least an acknowledgement that there is a need for an explanation of how the hearer sets the values of these parameters in particular cases. Obviously, a hearer will not come to know anything about the truth conditions for a particular utterance of (12a) unless he knows something like the conditional in (12b). However, if coming to know the truth conditions for an utterance of (12a) is coming to know the truth conditions of the thought that is communicated by it, then the hearer must also know whether and how the antecedent of the conditional is satisfied, for the representation of this thought must contain a representation of the female who the speaker is taken to be referring to. In other words, the truth conditions for thoughts are *not* determined by what the hearer knows about the semantic principles governing language.

As Carston has shown in a series of publications (see for example, 1988, 1993, 1997a, 1997b), the linguistic under-determination of the propositional content of utterances is not exhausted by referential indeterminacy. Apart from the problem raised by lexical ambiguity (for example (13)) and the unspecified scope of quantifiers (for example (14)), we must also consider how hearers are able to recover the intended propositional content from utterances with missing constituents. The range of examples is not restricted to utterances which contain inherently elliptical expressions or constructions (for example, (15–17), but also includes fully sentential utterances which are not generally classified as linguistically elliptical (for example (18)) and fragmentary utterances (for example (19–20)).

(13) The coach left the stadium at midday.
(14) Everyone has to go to the meeting.
(15) We went out for Christmas. The meal was really nice. But it wasn't
      the same somehow. [same as having Christmas at home]
(16) Don't sit on that rock. It'll fall. [if you sit on the rock]
(17) He works too hard. [for what?]
(18) Dogs must be carried. [if you are travelling on the London Underground with one]
(19) On the table.
(20) Lovely.

Moreover, there are utterances which, apart from their referential indeterminacy, have a meaning which determines a proposition with truth conditions, but not the proposition which is understood to have been expressed. For example, the hearer could recover a truth-evaluable proposition from (21) on the basis of its linguistic meaning and reference assignment. But this will not be a proposition which the speaker would have intended to communicate.

(21)   It'll take some time to get there.

It is not clear how one would capture the contribution made by linguistic meaning in these examples by means of conditional normal forms (cf. (12)). One would have to show that in each case there is a hidden indexical which imposes a constraint on the proposition expressed. While this may be feasible in cases of grammatical ellipsis, it is difficult to see how it would apply in cases such as (19), (20) or (21).[1] In other words, there are aspects of the truth conditional content of an utterance which cannot be determined by its linguistically determined semantic representation.

According to the standard versions of truth conditional semantics, the meaning of an expression is analysed in terms of the contribution it makes to the truth conditions of the sentence that contains it. However, when we shift our attention to the truth conditions of utterances, there is a further contextual entanglement not recognized by truth conditional semanticists. There are expressions which while they make a contribution to the truth conditions of the utterances which contain them, do not make the same contribution in every case. Thus as Carston (1997, 1998) shows, a speaker can use an expression to mean either something looser or something more restricted than the meaning it has taken out of context. For example, a child, angered by what he perceives as his mother's lack of maternal feeling, may produce the utterance in (22a):

(22)   (a) You're not a real mother.

On the other hand, the same child may well produce (22b) in order to clear up the confusion caused when his teacher addresses his childminder as his mother.

(22)   (b) She's not my mother, but only my childminder.

---

[1] See here the debate between Stainton (1994, 1997, forthcoming) who defends the idea that there are non-sentential assertions which have propositional content, and Stanley (2000, 2002) who argues that the phenomena Stainton calls non-sentential assertions are either grammatically elliptical and hence have a full sentential structure or are not genuine linguistic acts of communication at all and are more like winks or taps on the shoulder.

On the assumption that Higginbotham would give the conditional in (23b) as the linguistically determined semantic representation of (23a), it is difficult to see how he would capture the difference between the contribution made by *mother* to the truth conditions of (22a) and the contribution it makes to the truth conditions of (22b).

(23)   (a) She is a mother.
       (b) If *x* is referred to by *she* in the course of an utterance of (23a) and she is female, then that utterance will be true iff *adult female parent (x)*.

This chapter began with the question of whether truth conditional semantics could be regarded as a theory of semantic competence in the sense developed in the Chomskyan research programme. I have argued that a fundamental problem with grafting an externalist theory of truth conditional semantics on to an internalist theory of language is that the linguistic properties of an utterance do not fully determine a proposition with truth conditions – contextual entanglements are inevitable. This might be taken to suggest that alongside a theory of semantic competence we need a theory of pragmatic competence which would explain how we resolve these entanglements. In the next section I ask if there is any sense in which a notion of pragmatic competence is viable.

1.3   Pragmatic competence

Chomsky himself distinguishes *grammatical competence*, which he describes as the computational aspects of language that constitute knowledge of form and meaning, from *pragmatic competence*, which he defines as knowledge of the conditions for appropriate use, of how to use grammatical and conceptual resources to achieve certain ends or purposes (Chomsky 1980).

Unfortunately, Chomsky has said little further about the notion of pragmatic competence. But it is interesting to ask whether a system of knowledge about how to use language could solve the problems introduced in the last section. These are questions about how hearers recover the interpretations of utterances such as those in (12a) or (13–21) on the basis of the sometimes skeletal information provided by the grammar.

(12)   (a) She is lazy.
(13)   The coach left the stadium at midday.
(14)   Everyone has to go to the meeting.

(15)   We went out for Christmas. The meal was really nice. But it wasn't the same somehow. [same as having Christmas at home]

(16)   Don't sit on that rock. It'll fall. [if you sit on the rock]

(17)   He works too hard. [for what?]

(18)   Dogs must be carried. [if you are travelling on the London Underground with one]

(19)   On the table.

(20)   Lovely.

(21)   It'll take some time to get there.

A theory that explained how interpretations are recovered would have to actually integrate linguistic knowledge and non-linguistic knowledge. That is, it would have to provide a model of whatever does the interpretive work. It is difficult to see how this could be done by a system of knowledge.

It would seem then that a competence theory of pragmatic interpretation is impossible. And indeed Carston (1998, 1999) has argued that Kasher's (1991a, 1991b) attempts to develop Chomsky's brief remarks on pragmatic competence do not result in a theory of *competence* at all. Kasher (1991a,1991b) speaks of a number of different systems of pragmatic 'knowledge'. Of these it seems that only one is strictly language specific – the knowledge of basic speech act types, for example, assertions, questions and commands. The rest are domain-neutral or part of general knowledge. Thus in addition to a number of pragmatic 'modules', including one for basic speech acts and one for what he describes as 'talk-in-interaction', he proposes, first, a system called 'central pragmatics' which is defined as the knowledge of the general cognitive principles and general knowledge involved in the 'generation' of conversational implicatures, aspects of style and politeness, and, second, a system called 'interface pragmatics' which he defines as knowledge which is involved in integrating data from the language module and other sources (for example, in the assignment of reference to indexicals).

Carston (1998, 1999) argues that the principles which operate Kasher's central pragmatics are performance principles rather than principles which could be taken to constitute 'pragmatic knowledge' or 'competence'. For example, his general pragmatic principle which says 'Given a desired end, one is to choose that action which most effectively, and at least cost, attains that end' (1991b: 577) is a principle which guides behaviour. And his 'interface pragmatics', since it is a mechanism which integrates various inputs in accordance with a non-linguistic principle of best fit, must also be a performance mechanism.

For Carston this is not surprising for 'when it comes to internalist pragmatic theorising a shift from a competence to a performance perspective is virtually inevitable' (1999:93).

The performance perspective which Carston has in mind, and the one which informs this book, is not *linguistic* performance, but rather *cognitive* performance, or, more particularly, the inferential mechanisms which receive input from the linguistic performance mechanisms (the parser), perceptual sources (the senses) and conceptual sources (memory), and deliver interpretations (explicit and implicit content). In other words, according to this view, the point of contact between semantics and pragmatics is at the interface between the linguistic parser, which receives input from linguistic competence and delivers linguistically determined semantic representations, on the one hand, and the inferential mechanisms which take these semantic representations as input for the computations which deliver the representations which the hearer takes to be representations of the speaker's communicative intentions, on the other. These computations, like all performance mechanisms, are constrained by time and considerations of effort.

This is essentially the relevance theoretic picture of the relationship between linguistic form and pragmatic interpretation that I shall be outlining in chapter 3. Right now I wish to turn to a very different sort of case that has been made for a theory of pragmatic competence – a case which, as we shall see, has bearing on the sort of phenomena that are central to this book.

This case is made by Prince (1988). While she recognizes that there are aspects of interpretation, for example, conversational implicature, which fall outside the theory of linguistic competence, Prince draws attention to a range of phenomena which she argues are both pragmatic and linguistic. Consider, for example, the differences between the cleft structures in (24).

(24)   (a) It was Anna who played 'Summertime' at the Christmas concert.
       (b) It was 'Summertime' that Anna played at the Christmas concert.
       (c) It was at the Christmas concert that Anna played 'Summertime'.

These differences are pragmatic, according to Prince, because they are not differences in propositional content, but rather differences between the *contexts* in which the utterance of each sentence is felicitous. Thus (24a) but not (24b–d) is appropriate as an answer to (25a); (24b) but not (24a) or (24c–d) is appropriate as an answer to (25b); (24c) but not (24a–b) or (24d) is appropriate as an answer to either (25c) or (25d).

(25)   (a) Who played 'Summertime' at the Christmas concert?
       (b) What did Anna play at the Christmas concert?
       (c) At which concert did Anna play 'Summertime'?
       (d) Where did Anna play 'Summertime'?

Prince analyses these differences in terms of the distinction between focus and presupposition, so that 'the proposition conveyed is structured into two parts, one an open proposition [and the other] its instantiation' (1988:168). The felicitous use of one of these sentences 'requires that the open proposition be appropriately construed as shared knowledge' (1988:168). For example, the open proposition conveyed by (24a) is the one in (26).

(26)   x played 'Summertime' at the Christmas concert.

Prince suggests that this analysis of the pragmatic differences between these sentences could be unpacked along the lines suggested by Wilson and Sperber (1979). Thus the difference between (24a) and (24b) is that whereas the speaker of (24a) is communicating that (26) is the first 'background entailment' (Wilson and Sperber 1979), the speaker of (24b) would be communicating that (27) is the first background entailment.

(27)   Anna played x at the Christmas concert.

However, whereas Sperber and Wilson (1995) argue that the differences between (24a–c) are not linguistically determined, but the result of 'a natural linkage between linguistic structure and pragmatic effects' (1995:213), Prince argues that these pragmatic differences are arbitrary and language-specific and hence linguistically encoded. She argues against the view that they follow from an iconic property of *it*-cleft constructions by showing that the constructions used by other languages for performing the function performed by the English cleft may have 'a dramatically different' syntax (1988:168). For example, she compares English examples such as the ones in (28) with the corresponding Yiddish structures:

(28)   (a) ...zey hobn gefunen aykhmanen
          ...they have found Eichmann
       (b) ...*dos* hobn gefunen aykhmanen
          ...this they have found Eichmann
          ...it was they who found Eichmann
                              (Prince 1988: 169)

The English *it*-cleft involves syntactic subordination of the part of the sentence corresponding to the open proposition, and movement or isolation of the constituent corresponding to the instantiation of the variable in this proposition (together with marking of the variable's position, for example, with a trace). In contrast, in the Yiddish example, there is no subordination, no syntactic isolation of the focussed constituent and no trace. It consists of a single clause, has a post-verbal subject and has a dummy NP in first position that is not an argument of the verb. Moreover, as Prince points out, the *dos*-construction is syntactically equivalent to the *es*-construction in (29), which is used in a totally different context, namely, when the subject is non-thematic or, in other words, 'when the fewest assumptions about shared knowledge are warranted' (1988:169).

(29)   [Come to me, I've been away looking for you on twisted roads. I'm still young, inexperienced ...]
       (a) ...fremde mentshn kenen mikh farnarn
          ...strange people can me entice
          'Strange people can entice me'
       (b) ...*es* kenen fremde mentshn mikh farnarn (Shavaib: *Moyde ani*)
          ...it can strange people me entice
          'Strange people can entice me', 'It can happen that strange people entice me'
       (c) ?...*dos* kenen fremde mentshn mikh farnarn
          ...this can strange people me entice
          'It is strange people that can entice me'

Prince goes on to make the same kind of point about a range of other syntactic structures, for example, gapping, topicalization and VP preposing, as well as the referential options in examples such as (30).

(30)   (a) Last week I read a book and I met an author.
       (b) Last week I read a book and I met the author.

In each case she argues that cross-linguistic comparisons show that 'a significant part of a speaker-hearer's competence involves ... knowing which syntactic and referential forms trigger which nonlogical inferences' (1988:179).

The claim that these inferences are not due to the iconicity of the forms in question, while supported by cross-linguistic analysis, is controversial (see Sperber and Wilson 1995: 202–17). However, my aim here is not to enter into this controversy, but to locate Prince's work in the kind of theoretical