

RELEVANCE AND
LINGUISTIC MEANING

The semantics and pragmatics
of discourse markers

DIANE BLAKEMORE



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Cambridge University Press 2002

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2002

Printed in the United Kingdom at the University Press, Cambridge

Typeface Times 10/13 pt *System* L^AT_EX 2_ε [TB]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Blakemore, Diane.

Relevance and linguistic meaning : the semantics and pragmatics of discourse
markers / Diane Blakemore.

p. cm. – (Cambridge studies in linguistics; 99)

Includes bibliographical references and index.

ISBN 0 521 64007 5

1. Discourse markers. 2. Semantics. 3. Relevance. 4. Pragmatics.

I. Title. II. Series.

P302.35 .B58 2002

401'.41 – dc21 2002019248

ISBN 0 521 64007 5 hardback

Contents

| | | |
|----------|--------------------------------------|------------------|
| | <i>Acknowledgements</i> | <i>page viii</i> |
| | Introduction | 1 |
| 1 | Meaning and truth | 12 |
| 2 | Non-truth conditional meaning | 32 |
| 3 | Relevance and meaning | 59 |
| 4 | Procedural meaning | 89 |
| 5 | Relevance and discourse | 149 |
| | <i>Conclusion</i> | 184 |
| | <i>Bibliography</i> | 186 |
| | <i>Indexes</i> | 193 |

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 objects 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. Moreover, 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 consequent 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 farnam
 . . . strange people can me entice
 'Strange people can entice me'
- (b) . . . *es* kenen fremde mentshn mikh farnam (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 farnam
 . . . 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

picture that is being developed in this book. As interesting as this controversy may be, the question it raises about what we should do with linguistically encoded meaning which affects contextually determined aspects of interpretation must be addressed. For it seems clear that whatever we say about the nature of the phenomena discussed by Prince, there are linguistically encoded aspects of meaning which affect pragmatic interpretation – namely, many of the expressions which have been called discourse markers. For example, while it is uncontroversial that the differences between the following are due to the linguistically encoded meanings of *moreover* and *but*, it seems that these differences must be analysed in terms of the contexts in which they are appropriate. Note that both (a) and (b) are uttered by a single speaker.

- (31) (a) [I think you should accept this paper for the conference. It's well written and it's got wonderful examples.]
Moreover, it's right on the conference theme.
- (b) [I think you should accept this paper. It's well written and it's got wonderful examples.]
 ? *But* it's right on the conference theme.

In other words, knowledge of the meanings of these expressions is knowledge of the contexts in which the utterances that contain them are appropriate.

Since this is knowledge about the linguistic properties of particular linguistic expressions, it must be regarded as part of linguistic competence. Since this knowledge is knowledge about the meanings of these expressions, we can also say that it must be regarded as part of semantic competence. At the same time, since this knowledge must be analysed as knowledge about the contexts in which utterances are appropriate, or knowledge about discourse functions, it would seem that it is in some sense pragmatic knowledge. Prince (1988) resolves this tension by proposing a sub-component of linguistic competence called 'discourse competence' or 'pragmatic competence'. A theory of pragmatic competence is, for Prince, distinct from that theory of pragmatics whose domain includes such matters as the identification of context-dependent aspects of propositional content and the recovery of implicatures. On these matters, Prince agrees with the view outlined above: they involve inferences which are constrained by general cognitive principles, and hence lie outside the domain of the grammar.

So the question is whether we really need Prince's extra component of the grammar? In chapters 3 and 4, I shall show that within a relevance theoretic framework there is no need for such a component, and that the contribution of expressions such as *moreover* and *but* can be analysed within a theory of

linguistic semantics. However, as we shall see, the role of the semantic component is radically different from the one which I suspect Prince has assumed, for not only does it provide logical forms which are the input to those pragmatic processes which yield what is thought of as the propositional content of the utterance, but it also can provide information about the processes involved in using contextual assumptions for the recovery of the intended interpretation.

Such a semantic component would not, of course, be truth conditional, since it includes knowledge of the meanings of expressions and structures which do not play a role in the identification of truth conditions. The cleft constructions in (24a–c) and the utterances in (31a–b) are truth conditionally equivalent. However, as we have seen, the assumption that a sentence encodes a proposition with truth conditions cannot be maintained: there is a gap between what the sentence encodes and the interpretation that is recovered by the hearer. Moreover, there are aspects of what we think of as the propositional content of utterances in which the grammar seems to play no role at all (recall examples such as (19–21)). If linguistic semantics is not itself truth conditional, then the inclusion of so-called non-truth conditional meaning within its domain should not matter.

But can we abandon truth conditions like this? How can we divorce linguistic meaning from truth without abandoning the intuition that we use language to talk about the world?

1.4 Semantic retreat?

In fact, some theorists *have* abandoned this intuition, arguing that no expression of language should be analysed in terms of content. Thus although Anscombe and Ducrot's (1983, 1989) argumentation theory began simply as an attempt to accommodate non-truth conditional expressions (for example, *but*) by showing that meaning can have an argumentative component, it developed into a theory of radical argumentativism in which the function of language is argumentative rather than informative. As Iten (2000a, 2000b) has shown, while this approach has offered some important insights into the meanings of non-truth conditional expressions such as *but*,² it is not concerned with language or meaning as a cognitive phenomenon. Thus an utterance is not something produced by a particular speaker with particular intentions, but a collection of different points of view. And a point of view is not anchored in the mind of a particular person with particular thoughts. As Iten shows, this means that argumentation theory is unable to address the sort of questions raised in this chapter, namely, how

² For discussion of Anscombe and Ducrot's analysis of *but*, see chapter 4 below.

a hearer is able to recover the intended interpretation for context dependent expressions.

Anscombe and Ducrot's (1989) argument that language does not perform an informative function is based on the subjectivity of language. The fact that the meanings of some words involve subjective judgement is undeniable. And, as we have seen (in examples (22) and (23)), an expression can be used to make different contributions to the interpretation of an utterance in different contexts. However, as Iten (2000b) argues, this is not in itself a reason for saying that language is not used to represent the world *at all* – and, indeed Anscombe and Ducrot have admitted that there are some utterances which they cannot analyse in purely argumentative terms.

The arguments in the previous section were not arguments for abandoning the idea that language is used to describe the world or arguments for abandoning truth conditions altogether. The argument was simply that truth conditions cannot play a role within the semantic component of the grammar, or, in other words that the semantic representations generated by the grammar are sub-propositional and hence not bearers of truth values. Levinson (2000) describes this position as one of 'semantic retreat' (2000:240–2), arguing that a level of semantic representation which is not a representation of the world is extremely impoverished – so impoverished that it cannot even capture traditional sense relations (e.g. entailment, contradiction, hyponymy). Instead of pushing truth out of linguistic semantics, he recommends accommodating context dependence by letting pragmatics 'intrude into semantic representations and consequently into their interpretation' (2000:193) as, for example, in discourse representation theory (Kamp and Reyle 1993), or in file change semantics (Heim 1982). He tries to illuminate this point of view with the following metaphor:

In these proposals [Kamp and Reyle, Heim] there is a common slate, a level of propositional representation, upon which both semantics and pragmatics can write – the contribution may be distinguished by the colour of the ink: semantics in blue, pragmatics in red! Semantics and pragmatics remain modular 'pens' as it were: they are separate devices making distinctively different contributions to a common level of representation. The slate represents the semantic and pragmatic content of accumulated utterances, and it is this representation that as a whole is assigned a model theoretic interpretation. (Levinson 2000:193)

Unfortunately, this metaphor confuses rather than illuminates. Having made it clear that anyone who does not agree that semantics is all about the business of assigning model theoretic interpretations is guilty of 'semantic retreat', Levinson uses the word 'semantics' in the above quotation to refer to the

sub-propositional contribution made by linguistically determined meaning rather than the business of assigning model theoretic interpretations. That is, he seems to accept that we must distinguish between what the ‘semantics’ pen writes on the representation that is assigned truth conditions, on the one hand, and the assignment of truth conditions to the propositional slate where both pens have written, on the other. Levinson would, I believe, be right to distinguish the two enterprises. What the blue pen writes is one thing. Assigning truth conditions to a common propositional slate is quite another. However, he would have to say which one of these enterprises is semantics. He would also have to say something more about the common propositional slate. What kind of representation is it?

On the assumption that the grammar cannot have access to non-linguistic, contextual information, it is difficult to see how the propositional representation which is assigned truth conditions is a *linguistic* representation, as, I think, Levinson is arguing. For on this assumption, the red pragmatics pen cannot write on any level of representation in the grammar. But this is not to say that there is no mental representation at all that can be assigned truth conditions. Fodor (1998) has argued that thoughts are representations not just in the sense of being represented in the mind (in the same way as, say, grammatical representations), but also in the sense that their content is at least partly determined by their relationship with the external world. Indeed, he has argued that it is thoughts rather than English sentences which have truth conditions: ‘English has no semantics. Learning English isn’t learning a theory about what its sentences mean, it’s learning how to associate sentences with the corresponding thoughts’ (1998:9).

In other words, according to this view, there is no need to retreat from truth conditions *per se*. It is simply that the common propositional slate on which Levinson’s two pens write is not a system defined by the grammar, but a representation defined by the cognitive system that delivers the contents of our thoughts, beliefs, assumptions, desires and intentions. As we will see, some of the thoughts communicated by an utterance are communicated explicitly, while others are communicated implicitly. For example, the utterance in (32B) explicitly communicates the assumption that Anna has a meeting and implicitly communicates that she is not here.

- (32) A: Is Anna here?
B: She’s got a meeting.

On Fodor’s view both of these assumptions have truth conditions.

This is also the view of the relevance theoretic framework that I shall be arguing for in chapter 3. Thus Wilson and Sperber (1993) argue ‘the primary bearers

of truth conditions are not utterances but conceptual representations' (1993:23). However, although it is accepted that both explicitly communicated and implicitly communicated propositions have their own truth conditions, and 'are capable of being true or false in their own right' (Wilson and Sperber 1993:6), there is a tendency to single out a single explicitly communicated proposition – often referred to as the proposition expressed – as the one which carries the truth conditions of the utterance. The suggestion is that an utterance has truth conditions only in a derivative sense – that is, by virtue of expressing a proposition with truth conditions. As we shall see, within this framework the assumptions explicitly communicated by an utterance are not restricted to 'the proposition expressed'. So this raises the question of how we decide which of these propositions is the 'proposition expressed', and whether the notion of truth conditions need play any role in the answer.

As we shall see in the first section of chapter 2, a number of theorists have observed that our intuitions about the truth conditions of an utterance are variable and consequently unreliable. This observation will be borne out when we come to look more closely at so-called non-truth conditional phenomena. Bach (1999) and Neale (1999) have argued that the fact that our intuitions about the truth conditions of an utterance are so variable can be explained if we abandon the assumption that an utterance has a single set of truth conditions. This line of thought is, in fact, anticipated by Wilson and Sperber in the article cited above (1993). Thus although they assume throughout this paper that an utterance expresses a single proposition, they ask whether this approach is indeed justified, and speculate that the intuitions of an utterance as a whole are based on the assumption which makes the major contribution to relevance. As we shall see, their approach provides a cognitive framework in which these questions can be investigated. However, as Iten (2000b) has suggested, it also provides a framework for the exploration of the far more radical suggestion that there is no need, from a cognitive point of view, for the notion of 'the truth conditions of an utterance' at all.

This brings us back to Levinson's charge of 'semantic retreat'. As I hope I have made clear in this section, relevance theory is not retreating from semantics in the sense of retreating from truth conditions – not even if it adopts the radical view just described. Interpretations, according to relevance theory, are conceptual representations – conceptual representations of thoughts – and they have truth conditions. Nor is relevance theory retreating from semantics in the sense of retreating from the contribution of linguistically determined meaning (Levinson's blue pen). This contribution is, Levinson says, impoverished and schematic: it is a schema for the construction of a propositional representation

rather than a proposition. At the same time, however, linguistically determined meaning makes a contribution to utterance interpretation which would be excluded from semantics if it were defined in terms of truth conditions. The blue pen does not just write on the 'propositional slate'.

The approach that I have just introduced is embedded in the view that language is a vehicle for thoughts and desires. That is, it views language in cognitive terms. However, non-truth conditional meaning is often viewed as the property of a theory which views language not as a vehicle for thought but as a vehicle for action. Speech act theory grew out of a dissatisfaction with the methods of formal logic as applied to ordinary natural language, and in particular, with what was perceived as an over-emphasis on the descriptive function of language. A theory which claimed to provide an alternative to formal semantics by focussing on ordinary language use might perhaps seem to provide a natural home for all those expressions which could not be accommodated within truth conditional semantics. In the next chapter, I shall explore speech act theoretic approaches to non-truth conditional expressions and ask whether they have in fact led to a unitary account of non-truth conditionality.