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.