Contents

Preface ix
Author's preface xii

SECTION 1
Survey 1

1 The nature of language
In the beginning … 3
The design of language 4
Animal communication 8
Human language: endowment or accomplishment? 11
Language, mind, and social life 13

2 The scope of linguistics
Experience and explanation 17
Models and maps 18
Dimensions of idealization 20
Langue and *parole* 21
Competence and performance 24
Knowledge and ability 27

3 Principles and levels of analysis
Type and token 29
Principles of classification 30
Dimensions of analysis 32
Levels of analysis 35

4 Areas of enquiry: focus on form
The patterns of sound: phonetics and phonology 41
Sound segments 42
Syllables 43
Stress and intonation 43
The construction of words: morphology 45
Derivation and inflection 46
The combination of words: syntax 48
Grammatical systems 49
Constituent structure 50

5 Areas of enquiry: focus on meaning
Meaning in language: semantics 53
The meaning of words 53
Semantic components 55
Sense relations 57
Words and phrases 60
Meaning in context: pragmatics 61
Reference, force, and effect 62
Context and schema 63
Negotiation of meaning 65
Relations between utterances 66

6 Current issues
The scope of linguistics 69
The data of linguistics 72
The relevance of linguistics 75

SECTION 2
Readings 79

SECTION 3
References 115

SECTION 4
Glossary 125

Preface

Purpose

What justification might there be for a series of introductions to language study? After all, linguistics is already well served with introductory texts: expositions and explanations which are comprehensive and authoritative and excellent in their way. Generally speaking, however, their way is the essentially academic one of providing a detailed initiation into the discipline of linguistics, and they tend to be lengthy and technical: appropriately so, given their purpose. But they can be quite daunting to the novice. There is also a need for a more general and gradual introduction to language: transitional texts which will ease people into an understanding of complex ideas. This series of introductions is designed to serve this need.

Their purpose, therefore, is not to supplant but to support the more academically oriented introductions to linguistics: to prepare the conceptual ground. They are based on the belief that it is an advantage to have a broad map of the terrain sketched out before one considers its more specific features on a smaller scale, a general context in reference to which the detail makes sense. It is sometimes the case that students are introduced to detail without it being made clear what it is a detail *of*. Clearly, a general understanding of ideas is not sufficient: there needs to be closer scrutiny. But equally, close scrutiny can be myopic and meaningless unless it is related to the larger view. Indeed, it can be said that the precondition of more particular enquiry is an awareness of what, in general, the particulars are about. This series is designed to provide this large-scale view of different areas of language study. As such it can serve as a preliminary to (and precondition for) the

models will work to different scales and give preference to different features. Like maps, all models are simplified and selective. They are idealized versions of reality, designed to reveal certain things by concealing others. There can be no all-purpose model, any more than there can be an all-purpose map. Their validity is always relative, never absolute. They are designed to explain experience, and so they should not be expected to correspond with it. None of them can capture the truth, the whole truth, and nothing but the truth. If they did that, they would cease to be models, of course, just as a map which corresponded exactly to the terrain would cease to be a map. In both cartography and linguistics the problem is to know what scale to use, what dimensions to identify, and where, in the interests of explanation, to draw the line between idealized abstractions and actual particulars.

Dimensions of idealization

If we consider the actual particulars of language, they appear to be a bewildering assortment of different facets. As a means of interaction between people, language is a social phenomenon. It enables us to give public expression to private experience and so to communicate and commune with others, to arrive at agreed meanings and to regulate relationships. For this purpose to be served, different languages have to be relatively stable codes which people contract into as a condition of membership of the communities that use them, and there have to be generally agreed ways of using the language in different kinds of social context. In this sense, to learn a language is an act of social conformity.

At the same time, language provides the means for non-conformist self-expression as well. There is always some room for individual manoeuvre. For example, an individual speaking French, or Swahili, or Chinese in the natural course of events will on the one hand produce instances of that language, combinations of words, in accordance with the underlying systems of rules and established meanings which constitute the linguistic codes in each case. But on the other hand, they will be producing unique expressions in the language by exploiting the potential of the code. Although individuals are constrained by conventions of the code and its use, they exploit the potential differently on different

occasions and for different purposes. But this conscious exploitation is not the only source of variation. The patterning of a person's use of language is as naturally distinctive as a fingerprint. And even spoken utterances repeated by the same person, though they may sound identical, are never acoustically alike in every particular. It is obviously socially necessary to assume that certain things are the same, even if, on closer scrutiny, they turn out to be different.

The point then is that, from one perspective, language is a very general and abstract phenomenon. It is a shared and stable body of knowledge of linguistic forms and their function which is established by convention in a community. At the same time, it is very particular and variable if we look at the actuality of linguistic behaviour. Since social control is necessarily a condition on individual creativity, there is no contradiction here. It is simply that the nearer you get to actuality along the scale of idealization, the more differences you discern as the more general abstractions disappear. It is therefore convenient to mark off limiting points along this scale to define the scope of linguistic enquiry.

Langue and *parole*

One such mark was made by Ferdinand de Saussure, the Swiss scholar usually credited with establishing the principles of modern linguistics. In a celebrated series of lectures in the early part of the century, he proposed that linguistics should concern itself with the shared social code, the abstract system, which he called ***langue***, leaving aside the particular actualities of individual utterance, which he called ***parole***. *Langue* was, on his account, a collective body of knowledge, a kind of common reference manual, copies of which were acquired by all members of a community of speakers. This distinction between language as abstract system and actual speech can be justified on two grounds (and it is not always entirely clear which one Saussure is arguing for). Firstly, it is convenient in that it delimits an area of enquiry which is manageable: it is possible in principle to conceive of a linguistics of *parole*, but the individual particularities of actual acts of speech are so varied and heterogeneous as to be elusive of description. Secondly, the concept of *langue* can be said to capture the central and determining aspect of

language itself. On this account, *parole* is the contingent executive side of things, the relatively superficial behavioural reflexes of knowledge. So *langue* can either be seen as a convenient principle of linguistics, or as an essential principle of language itself, or both.

There are a number of issues arising from Saussure's distinction. To begin with, one should note that the concept of *langue* eliminates from language its intrinsic instability. Language is necessarily, and essentially, dynamic. It is a process, not a state, and changes over time to accommodate the needs of its users. In fact Saussure was well aware of this. He was himself schooled in the tradition of historical linguistics which sought to account for changes in language over time, its **diachronic** dimension. But he conceives of *langue* as a cross-section of this process at a particular time, a **synchronic** state, which might be represented in the following diagram:

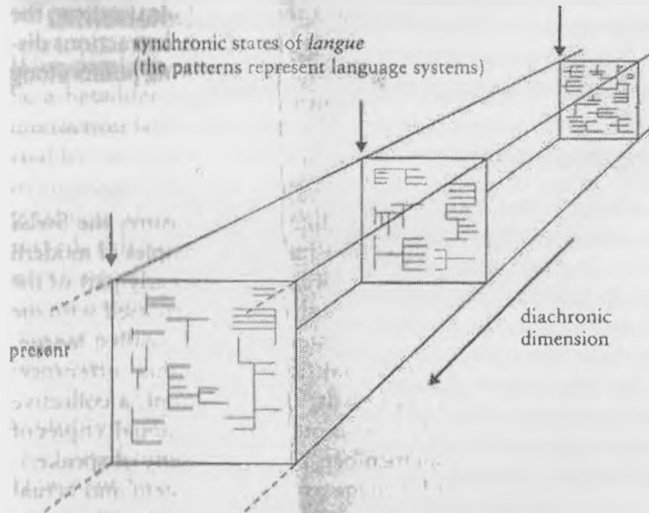

FIGURE 2.2 *The relationship between synchronic and diachronic aspects of language*

One difficulty about this conception, however, is that there is a confusion between synchrony and stability. Wherever you take a synchronic slice through language you will find not fixity, but

flux. This is because language does not just change *over time* but varies *at any one time*, and indeed this cannot be otherwise because the members of a community which 'shares' a language will themselves be of different ages, will use language differently, and will have different communicative and communal uses for it. Different generations generate differences. No matter how small the period of time, or limited the variety of language, there will be variations within it as it is fine-tuned by the community of its users. And as some of these variable uses become conventionalized, so they become established as changed forms. In other words, diachronic change over time is simply, and inevitably, a result of synchronic variation at any one time.

To illustrate his synchrony–diachrony distinction, Saussure drew an analogy with the game of chess. The synchronic cross-section of language (the state of *langue*) is, he argued, like the state of play at one time. We can study the disposition of the pieces on the board without considering the diachronic dimension of the game, that is to say, the moves that were made beforehand, or those that might be planned in the future. We can, in other words, see the pattern of pieces as a state of play and disregard it as a stage in the game. The analogy breaks down, however, because of course the game of chess is of its nature a sequence of separate stages and the game itself stops as each player takes a turn. But language is a continuity with no divisions of this kind. It is linguistics which makes it stop.

To say that diachrony and synchrony are not in reality distinct dimensions is not to invalidate the idealization that makes them distinct, but only to set limits on its claims to absolute validity. And this, as has been pointed out, is true of *all* models of language. If we wished to account for variation and change, we would draw the lines of idealization differently, but there would still be idealization. And the resulting model would necessarily be less revealing of the relative stability of language which serves as the necessary frame of reference in accounting for variation. You have to assume fixed points somewhere as bearings on description.

And as bearings on behaviour. It is important to note too that this assumption of stability can have a reality of its own. It is not only Saussure who conceives of language as a stable state. Although a close scrutiny of an actually occurring language will

reveal all manner of variation, people in the communities who speak it might well nevertheless *think* of their language as being settled and established, and accept the validity of grammars and dictionaries which record it as such. Members of a linguistic community may not have identical copies of *langue* in their heads, but they may nevertheless *believe* they do, and may consider whatever differences they do discern as matters of no real significance.

Competence and performance

A comparable distinction to that of Saussure, designed to idealize language data, and to define the scope of linguistic enquiry, is made by Noam Chomsky. He distinguishes **competence**, the knowledge that native speakers have of their language as a system of abstract formal relations, and **performance**, their actual behaviour. Although performance must clearly be projected from competence, and therefore be referable to it, it does not *correspond* to it in any direct way. As with other aspects of human life, we do not necessarily act upon what we know, quite simply because actions are inevitably caught up in particular circumstances which set constraints and conditions on what we do. So it is that actual linguistic behaviour is conditioned by all manner of factors other than a knowledge of language as such, and these factors are, according to Chomsky, incidental, and irrelevant to linguistic description. Performance is particular, variable, dependent on circumstances. It may offer evidence of competence, but it is *circumstantial* evidence and not to be relied on. Abstract concepts of competence and actual acts of performance are quite different phenomena and you cannot directly infer one from the other. What we know cannot be equated with what we do.

Chomsky's distinction obviously corresponds in some degree to that of Saussure. It represents a similar dichotomy of knowledge and behaviour and a similar demarcation of the scope of linguistic enquiry. There are, however, differences. To begin with, there is no ambivalence in Chomsky as to the status of the distinction. It is not that competence is presented as a *convenient* construct and therefore a useful principle for language study: it is presented as a *valid* construct, as the central principle of language itself. To focus on competence is to focus on what is essential and

primary. Performance is the residual category of secondary phenomena, incidental, and peripheral.

A second point to be made is that though *langue* and competence can both be glossed in terms of abstract knowledge, the nature of knowledge is conceived of in very different ways. Saussure thinks of it as socially shared, common knowledge: his image is of *langue* as a book, printed in multiple copies and distributed throughout a community. It constitutes, therefore, a generality of highest common factors. But for Chomsky competence is not a social but a psychological phenomenon, not so much printed as *imprinted*, not a shared generality but a genetic endowment in each individual. Of course, individuals are not innately programmed to acquire competence in any particular language, but competence in any one language can nevertheless be taken as a variant in respect to universal features of language.

Langue, then, is conceived of as knowledge which is determined by membership of a social community, and so it follows that the focus of attention will naturally be on what makes each *langue* different. In this definition of linguistic knowledge, the main question of interest is: what is distinctive about particular *languages* as social phenomena? Competence, on the other hand, is conceived of as knowledge which is determined by membership of the human species and it follows that the interest here will naturally be not on what makes individual competences different but what makes them alike. In this definition of knowledge the main question of interest is: what is distinctive about *language* in general, and as specific to the human species?

Chomsky's distinction, then, leads to a definition of linguistics as principally concerned with the universals of the human mind. Indeed, he has defined linguistics as a branch of cognitive psychology. His idealization is a strictly **formalist** one in that it fixes on the forms of languages as evidence of these universals without regard to how these forms function in the business of communication and the conduct of social life in different communities. In this respect, Chomsky's definition of competence as the proper concern of linguistics is much further along the continuum of abstraction than is Saussure's definition of *langue*, in that it leaves social considerations out of account entirely.

Two further issues are perhaps worth noting in respect to this

formalist definition of language. First, as was indicated earlier, it is obvious that the further one proceeds in abstraction, the greater the risk of losing contact with the actuality of language in use. If competence is knowledge of the abstract principles of linguistic organization, which may not be evident in actual behaviour, nor even accessible to consciousness, then what, one might reasonably ask, counts as empirical evidence for its existence? The answer to this question has generally been that linguists themselves, as representative native speakers of a language, can draw evidence from their own intuitions. But there seems no reason why one should suppose it as self-evident that linguists are reliable informants: on the contrary, one might more reasonably suppose that as interested parties with an analytic bent they would on the face of it be very untypical, and so be disqualified as representative speakers. There are ways of countering this argument, but problems about the link between abstraction and actuality remain, and the further language is removed from its natural surroundings, the greater the problem becomes. On the other hand, the more you locate it in its natural surroundings, the less you see in the way of significant generalization. The dilemma of idealization we discussed earlier will always be with us.

Whereas this first issue has to do with the methodology of linguistic enquiry, with how to give support to the statements you make, the second has to do with the scope of linguistic enquiry, with what your statements should actually be *about*.

And here we find something of an apparent paradox in Chomsky's position. What he represents as central in language is an abstract set of organizing principles which both define an area of human cognition, a specific language faculty, and determine the parameters of Universal Grammar. The various forms of different languages are of interest to the extent that they can be seen as alternative settings for these general parameters. The communicative functions such forms take on in actual contexts of use are of no interest at all. They furnish no reliable evidence of underlying cognitive principles: there are too many distractions in the data by way of performance variables. So the most important thing about language from this point of view is that it is evidence for something else, namely a faculty in the human mind, uniquely and innately specific to the species. In a sense, therefore, it would appear that what is

central in language is that it is not of itself central. Paradoxically, for Chomsky, the study of language depends on disregarding most of it as irrelevant. Indeed, in this view, what linguistics is about is not really language but grammar, and more particularly that area of grammar which is concerned with the structural relations of sentence constituents, that is to say, with **syntax**.

Chomsky's specification of the scope of linguistics is extremely broad and far-reaching in respect to its implications, encompassing as it does nothing less than the universals of the human mind. But it is, of course, correspondingly extremely narrow and inward-looking in respect to the familiar phenomenon of language itself. What Chomsky presents is an abstract explanation of language which is a long way from actual experience. Not surprisingly, it has been challenged.

Knowledge and ability

One objection to Chomsky's model is that it defines the nature of linguistic knowledge too narrowly to mean a knowledge of grammatical form, and more specifically of syntax. Knowing a language, it is objected, involves more than knowing what form it takes: it involves knowing how it functions too. And this in turn implies knowing about words, not just as formal items, constituents of sentences, but as units of meaning which interact with syntax in complex ways. The formal systems of a language, after all, have evolved in association with words as the internal semantic encoding of some external social reality. So an account of grammatical knowledge, the argument runs, cannot ignore the fact that linguistic form is functionally motivated, so that to abstract form so completely from function is to misrepresent the nature of language. In this view, linguistics is essentially the study of how languages *mean*, how they are functionally informed: it is **semantics** which is primary.

Chomsky's *formal grammar* seeks to identify particular features of syntax with reference to universal and innate principles of human cognition. An alternative is to think in terms of a *functional grammar*, to consider how language is differentially influenced by the environment, how it is shaped by social use, and reflects the functions it has evolved to serve.

But it is also argued that knowing a language also includes knowing how to access grammar, and other formal features of language, to express meanings appropriate to the different contexts in which communication takes place. This too is a matter of function, but in a different sense. Here, we are concerned not with what the language means, that is to say, the *internal* function of forms in the language code, but with what people mean *by* the language, that is to say, what *external* function forms are used for in communication. Knowledge in the abstract has to be made actual and this is normally done by putting it to communicative use, not citing random sentences. People do not simply display what they know. They act upon it, and their actions are regulated by conventions of different kinds. So, according to this point of view, competence is not only knowledge in the abstract, but also ability to put knowledge to use according to convention.

There are then two ways of revising Chomsky's conception of competence, of redrawing the lines of idealization in devising a model of language. Firstly, we can redefine what constitutes the code or internal language by including aspects which reflect the nature of language as a communicative resource. This results in a functional grammar and, we may say, broadens the concept of linguistic *knowledge*.

Secondly, we might extend the notion of competence itself to include both *knowledge* and the *ability* to act upon it. Performance, then, becomes particular instances of behaviour which result from the exercise of ability and are not simply the reflexes of knowledge. Ability is the executive branch of competence, so to speak, and enables us to achieve meaning by putting our knowledge to work. If we did not have this accessing ability, it can be argued, the abstract structures of knowledge—this purely *linguistic* competence—would remain internalized in the mind and never see the light of day. We would spend all our lives buried in thought in a paralysis of cognition. Since this ability is only activated by some communicative purpose or other, we can reasonably call this more comprehensive concept **communicative competence**.

3
Principles and levels of analysis

However linguistic knowledge is defined, it involves an abstraction from actuality, some kind of classification of experience. To say that you know a language implies that you have *inferred* certain generalities from particulars. That is what we do in language learning. To say you know how to act upon your knowledge implies that you can reverse the process and identify instances, that is to say, *refer* particulars to generalities. That is what we do in language use.

Type and token

It follows that linguistic description deals in generalities, in abstract **types** of language element of which particular instances are actual **tokens**. Consider, for example, the following line from Shakespeare's *Richard II*:

I wasted time and now doth time waste me.

On one count, there are nine word elements here, and thirty-two letter elements. This is a count of token occurrences. But the word 'time' occurs twice, so if we count word types, there are eight words here. Similarly, if we count letter types, there are ten, since the letters 'i' and 'w' occur three times, 't' five times, and so on. But if we define elements differently, we would, of course, get other counts. Thus, we might count 'wasted' and 'waste' as tokens of the same type (the verb 'waste') or as different types if we are thinking in terms of **lexical items**, since the verb is used in two different meanings, 'to use extravagantly' and 'to make weaker and thinner'. Or we could adjust our focus again and consider vowels and

'Would you like to have your lunch now?'; 'You have left my door open, and I would like you to close it.' In other words, the formal properties of language are functionally motivated.

Within the noun phrase, then, there are tight structural constraints on sequence. The noun phrase itself, as a higher level constituent, is allowed more room for manoeuvre. We saw this earlier (in Chapter 3) when we noted the different ordering of constituents in the sentence: 'People ride bikes in Oxford'/'In Oxford, people ride bikes'/'People, in Oxford, ride bikes.'

But just as the tightness of control within the noun phrase is motivated, so is this relaxation of control of constituents within the sentence. Generally speaking, the larger the constituent, the greater its mobility. In all cases, the syntax provides a means to exploit more fully the meanings that are encoded in words.

The principles of constituent structure, based as they are on the syntagmatic and paradigmatic relations discussed earlier (in Chapter 3) are very powerful. They can produce (or generate) elaborate combinations and permutations of all kinds. Linguists will often demonstrate this by the invention of sentences of curious and baffling complexity, sentences which bear little resemblance to what people actually produce as utterances in real life. We have to bear in mind, however, that these sentences are devices for illustrating the syntactic means which speakers of a language have at their disposal, not the ways in which they actually employ them in contexts of use.

The morphological and syntactic processes which have been briefly outlined can be described in purely formal terms as operations of the code. But it is important to recognize that they function as devices for extending word meanings, and so constitute a communicative resource.

5 Areas of enquiry: focus on meaning

Meaning in language: semantics

The study of how meaning is encoded in a language is the central business of semantics, and it is generally assumed that its main concern is with the meanings of words as lexical items. But we should note that it is not *only* concerned with words as such. As we have seen, meaning also figures at levels of language below the word and above it. Morphemes are meaningful, for example: the derivational prefix *pre*- means 'before', so a '*pre*^fix' means 'something fixed before'. '*Un*^fixed' means 'not fixed', '*re*^fixed', 'fixed again'. The inflectional morphemes are meaningful too: 'fix^*ed*' signals 'past' in contrast with 'fix^*es*' which signals 'present' (and third person subject as well). Semantics is also necessarily implicated in syntax. As we saw in Chapter 3, the constituent structure 'People in Oxford/ride/bikes' means something different from 'People/ride/bikes/in Oxford'. Similarly, 'The bishop offended the actress' and 'The actress offended the bishop' are quite distinct in meaning, because word order is a syntactic device in English and so we assign subject status to the first noun phrase in each case. In both examples we have exactly the same collection of words; it is only the way they are ordered that makes them different.

The meaning of words

Facts like these have sometimes led linguists to undervalue the significance of the lexical meaning of words. It is common practice to expose the semantic indeterminacy of words in juxtaposition by citing ambiguous newspaper headlines like:

SQUAD HELP DOG BITE VICTIM
ASIANS SETTLE IN WELL

The words alone will not do, it is argued: only grammar can sort out the ambiguity by identifying different constituent structures ('settle in/well' vs 'settle/in well', for example). And the argument is often further illustrated by quoting from Lewis Carroll's 'Jabberwocky' to show the superior semantic signalling of grammar: For example,

'Twas brillig and the slithy toves
Did gyre and gimble in the wabe.

The words, it is claimed, are nonsense and so all we can do is identify the form classes on grammatical evidence: adjectives 'brillig' and 'slithy', verbs 'gyre' and 'gimble', nouns 'toves' and 'wabe'. So it is that whatever meaning can be gleaned from these lines must depend entirely on the grammar. But this does not seem to be so. Although these words are not part of the normal vocabulary of English, they resemble words that are, and so we treat them as lexical items and assign them meaning accordingly. Thus, 'brillig' can be said to suggest 'brilliant/bright', 'slithy', 'slimy/lithe', and 'wabe', 'wave'. So for me, at least, these lines project some meaning roughly on the lines of: 'It was a bright day, and reptilian creatures were frolicking in the waves'. Other people will no doubt read the lines differently, but they will do so by assigning some meaning or other to the lexical items. They will not just ignore them. Meaning may not be fully determined by lexis, but given a collection of words, as we saw with the artist and the church in Chapter 4, we can always infer *some* figment of a proposition. Grammar actually provides much less to go on. Nobody, I imagine, would make much sense of:

'Twas adjective and the adjective nouns
did verb and verb in the noun.

So although meaning is indeed signalled, as we have seen, by the morphological and syntactic processes of word adaptation and assembly, this is far from the whole story. Obviously enough these processes need words to work on, and it is the words which provide the main semantic content which is to be selected from and shaped. The grammatical processes we have discussed can be

seen as playing a supportive role whereby existing units of lexical meaning are organized, modified, and tailored to requirements. They do not initiate meaning; they act upon meaning already lexically provided.

Semantic components

What kinds of meaning, then, are encoded in the word? We can begin by referring to the same principle of constituent assembly that has served us so well so far. When considering inflectional affixes in the last chapter, it was pointed out that it is common to find two morphemes fused into one form, as in 'come' + past tense = 'came'. When considering derivational affixes we noted that '*un–*' and '*re–*' can combine with various lexical items to yield others like 'unfix', 'undo', 'unscrew', 'refix', 'retell', and 'review', and so on. We have already established the semantic character of these morphological forms. We can say, then, that a lexical item like 'unfix' has two semantic elements or components, each given separate expression in the word form '*un* + fix'.

Now it happens that many such derived forms have semantic equivalents which are single morphemes: 'unwell' = 'sick', for example, 'unhappy' = 'sad'. Furthermore, there are many equivalents which can take the form not of single words but of phrases where the bound morpheme separates itself from bondage and becomes free. So 'unwell' = 'not well', 'unhappy' = 'not happy', 'reborn' = 'born again', 'replant' = 'plant again', and so on. In George Orwell's novel *1984*, this principle of decomposition provides the basis for the reformed English of Newspeak: in Newspeak, for example, 'excellent' becomes 'plusgood', 'bad' becomes 'ungood', 'terrible' becomes 'plusungood', and so on.

Now (without commitment to the principles of Newspeak) it seems reasonable to suggest that a lexical item like 'sick' is a version of 'unwell': it is just that the two morphological elements have become fused into one. It would follow that if 'unwell' has two elements of meaning or **semantic components**, then so does 'sick'. And if these lexical items can be said to be encodings of different semantic components, then it would also seem logical to suppose that the same can be said of all lexical items, the only difference being that such components are explicitly signalled in some cases, but not in others.

The signalling is not straightforward, however. When a free lexical form becomes bound as an affix, its meaning is not just added, but acts upon the host lexical item in various ways. Thus, 'careful' can be analysed as 'full of care', but 'careless' does not mean 'with less care' but 'with *no* care'. Some affixes activate grammatical relations. The suffix '*–able*', for example, contracts a passive relation with its stem. So 'eatable', for example, means not 'able to eat' but 'able *to be eaten*'. With '*–less*' and '*–able*', the semantic effect of affixation is predictable. In other cases things are not so simple. The suffix '*–er*' derives a noun from a verb, and so denotes an actor of the action. Thus, words like 'baker' and 'keeper', can be taken apart and glossed as 'a person who bakes' and 'a person who keeps (something)'. Here the actor is a human agent. But it can also be an inanimate instrument. A 'cooker' is not 'a person who cooks' but 'a device for cooking', and in words like 'printer', 'cleaner', and 'speaker' the suffix can denote either agent or instrument. And with words like 'creeper' (meaning 'plant') and 'breaker'(meaning 'wave') the original significance of the suffix has now, in part at least, disappeared.

And it is commonly the case, of course, that the distinctive meaning of the lexical host disappears and blends in with the affix in the historical process of etymological change. So it is with words like 'reckless' or 'feckless', which cannot mean 'with no reck' or ' with no feck' since there are no such lexemes in English. Conversely, when an affix attaches itself to an existing form, it may blend with its host, and again the lexical whole is not a sum of its parts. The prefix '*re–*' is even more unreliable in this respect than the suffix '*–er*' referred to above. 'Return' does not normally mean 'turn again' or 'recall' 'call again'. When they do signal such meanings, they are generally given a hyphen in writing and marked stress in speech to indicate that the prefix retains its semantic identity. Thus, you have a 're-call' ('ri:kɔ:l) button on the telephone, but you may not 'recall' (rɪ'kɔ:l) how to use it.

The general point is, then, that we can conceive of all lexical items as encodings of one or more semantic elements or components, whether these are overtly signalled or not, and in identifying them we can establish the **denotation** of words. Thus, one denotation of the verb 'return' can be specified as [come + back], another as [give + back]. 'Come/go' and 'give/take' in turn can be

said to consist of components: something along the lines of [move + self + towards/away] on the one hand, and [move + something + towards/away] on the other.

These components of meaning can be seen as analogous with segments of sound, as discussed in Chapter 4. The same principle of combination is at work. In our previous discussion, we were able to establish contrasts between phonological words by invoking minimal differences in the sequence of sound segments. Thus, 'come' (/kʌm/) contrasts with 'gum' (/gʌm/) with respect to the one feature of voice on the initial consonant—i.e. the sounds /k/ and /g/ are formed in exactly the same way, except that in /g/ the vocal cords vibrate and in /k/ they do not. The same principle applies here: we can establish similar minimal pairs of lexical items with respect to their semantic components. Thus, 'come' contrasts with 'go' in respect to the one feature of directionality: [movement + *here*] as opposed to [movement + *there*].

This approach, known as **componential analysis** thus provides an inventory of the semantic features encoded in lexical forms. It can, of course, become immensely complicated and unwieldy, and as in all analysis, as the details proliferate they can lose their point and create confusion. The essential purpose of componential analysis is to identify certain general conceptual categories or semantic principles which find expression in the particular components. Among such categories are state, process, causality, class membership, possession, dimension, location, and, as we have seen with 'come' and 'go', directionality. By invoking them, we can move on from the denotation of particular lexemes to the **sense relations** that exist between them.

Sense relations

Consider directionality, for example. As we have seen, it provides the basis for the distinction between 'come' and 'go'. But it also figures in other contrasts as well, for example, 'give/take', 'advance/retreat', 'arrive/depart', 'push/pull', 'send/receive', and 'buy/sell'. All of these pairs have the common feature of process, but the terms in each pair express opposite directionality, and in this respect are examples of **antonymy**. And within this group, we can distinguish a sub-set of which 'give/take' and 'buy/sell' are members. Here, there is a relation of reciprocal implication,

known as **converseness**: 'sell' necessarily implies 'buy' and vice versa (if X sells a car to Y, Y necessarily buys the car from X). However, this sense relation is independent of directionality. Not only does it exist between the locational terms 'above/below', for example (if A is above B, B is necessarily below A), but also between such reciprocal roles as 'parent/child', where the sense and family relations, so to speak, coincide: 'If Anne is Harry's child, he is her parent'.

If we now consider a different semantic feature, that of dimension, we come to a meaning opposition (or antonym) of a rather different kind. Consider the adjective pairs: 'big/small', 'long/short', 'thin/fat', and 'far/near'. Here, we have not absolute but relative oppositeness: not either/or but degrees of difference in respect to some norm or other. Thus, a large mouse is a small animal as compared with a small elephant—or even a very small elephant—which is a large animal. Adjectives of this kind are said to be *gradable*. They can, naturally, occur with intensifiers (for example, 'very', 'extremely') and with comparative and superlative degrees (for example, 'smaller', 'smallest'). Again, as with the directional component above, this kind of antonymy is by no means restricted to lexical items with a dimensional component. 'Hot/cold', 'old/new', and 'happy/unhappy' are gradable, for example. 'Male/female', and 'married/unmarried', on the other hand, are not. You can be 'very happy' or 'rather old' but not (normally) 'rather female' or 'very married'.

The examples 'happy/unhappy', and 'married/unmarried' bring us to another sense relation. According to the earlier argument, these items with their explicit prefixes '*un–*' are equivalent in denotation to fused versions 'unhappy' = 'sad', 'unmarried' = 'single'. With the prefixed versions, the antonymy is explicitly signalled. But there are innumerable other examples where two lexical items will contract exactly the same opposition: 'buy/sell' = 'purchase/sell', 'arrive/depart' = 'arrive/leave', and so on. To the extent that 'buy' and 'purchase', and 'depart' and 'leave' are relational equivalents, they can be said to be examples of **synonymy**.

Earlier we analysed 'come' as consisting of the features [move + towards]. But 'move' as a semantic feature figures in the denotation of countless other lexical items as well of course. Thus, 'walk' is 'to move on foot'. But 'walk', too, is semantically incorporated

into other words: 'march', 'amble', 'stroll', 'tramp', and 'stride', for example. 'Walk', then, is the general or **superordinate** term, and the others, the more particular instances included within it, are its subordinate terms or hyponyms. In the same way, 'animal' is a superordinate term, 'mouse' and 'elephant' are hyponyms. But we can establish intervening levels of **hyponymy**: 'mouse' is a hyponym of the superordinate 'rodent' (together with the co-hyponyms 'rat', 'porcupine', etc.), while 'rodent' is a hyponym of the superordinate 'mammal', which is in turn a hyponym of 'animal'.

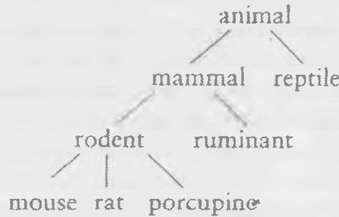

FIGURE 5.1 *Part of a hyponymic tree for 'animal'*

Each superordinate necessarily possesses a semantic feature common to all its hyponyms. To the extent that each co-hyponym has a distinct semantic specification, it serves as a superordinate to the next level of classification down, until all distinctive features are exhausted. It follows that where two lexical items appear in the same position on the tree as hyponyms we have synonymy. We may decide, for example, that 'amble' and 'stroll' are not distinguishable as ways of walking, and so are synonyms in that they have the same hyponymic relation to the superordinate word 'walk'. Notice, though, that this has to do with the equivalence of denotation as elements of the code. Synonymy as discussed here is a *semantic* relation. The extent to which synonyms have a different range of functions when they are actually put to use in contexts of communication is a different matter, which we will be taking up a little later in this chapter in the discussion of pragmatics, or meaning in context.

words and phrases

We began this chapter by looking at ways in which semantic components are overtly signalled by derivational affixes as parts of words, and we have subsequently considered how words themselves as lexical items relate semantically to others. Lexical items, however, do not only come in the form of single words. They appear as pairs, for example, in phrasal verbs ('see to', 'look up', 'pass by') or compound nouns ('prime minister', 'postage stamp', 'table lamp').

But lexical items come as larger clusters of words as well. Take the single word 'often' (morphologically simple) and the word 'repeatedly' (morphologically complex with its affixes). These can be seen as synonymous with the expressions 'over and over again' or 'time after time' respectively. Such expressions are **formulaic phrases**, and since they are complete units of meaning semantically equivalent to single words, they too can be considered as single lexical items. What distinguishes them is that in their case it happens that the semantic elements have not fused into a single form but find expression as separate words in a composite unit.

But it needs to be noticed that these lexical phrases are *compounds* of words, and not, as with the syntactic phrases that were considered in the previous section, *combinations* of words. Thus, the words in the expression 'time after time' are separate, but they are not independent as grammatical constituents. So we would not, for example, treat 'time' as a normal noun and pluralize it (*'times after times'), or add an article (*'*a* time after *the* time'), or replace 'after' with another preposition (*'time *before* time'). The words are compounded into a fixed association which syntax cannot meddle with. There are innumerable instances of such compounded lexical items in English, as there are in other languages: 'many a time and oft', 'by hook or by crook', 'easy come easy go', 'easier said than done', 'run of the mill', 'in the main', 'by and large', 'least said, soonest mended', and so on.

So some sequences of lexical items, or **collocations** are fixed, but there are innumerable others which are not and which *can* be syntactically modified. But only up to a point. Here, we come to the uncertain border between lexis and syntax, where words move from a compounding to a combining relationship.

Take the common expression 'He thought better of it'. Here, the subject is a normal sentence constituent and so can be replaced by an infinite number of other noun phrases ('I', 'You', 'They', 'The retired generals', 'The poor old pensioner living next door ...'). But although the noun phrase thus combines freely with what follows, the rest of the expression is resistant to replacement. It would be odd English to say: *'He reflected better of it', *'I thought worse of it', *'They thought better about it'. So this expression 'He thought better of it' consists in part of constituents which are combined and in part of lexical items which are compounded. It is not entirely fixed, as is 'by hook or by crook', for example, which is a complete lexical compound in that it admits of no interference at all (*'by the hook and the crook', *'by hooks or crooks', etc.). But it is not entirely free either, like 'He thought about it' which consists of a straightforward combination of sentence constituents. Grammatical rules can be seen as devices for regulating the meaning of words. The difficulty is that they are not completely regular in their application.

All this may seem to be fairly trivial—a detail or two about the peculiarities of English. But it illustrates again that semantics is not only a matter of assigning meaning to individual units, whether these be morphemes or words, but is also concerned with the *relationships* between them, how they act upon each other, how they fuse, compound, and combine in different ways. Semantics is the complex interplay of morphology, lexis, and syntax. Complex though it is, however, it does not account for all aspects of meaning. We still have pragmatics to consider.

Meaning in context: pragmatics

Semantics is the study of meaning *in* language. It is concerned with what language means. This is not the same as what people mean *by* the language they use, how they actualize its meaning potential as a communicative resource. This is the concern of **pragmatics**.

The distinction is easy enough to demonstrate. Consider the expression:

The parson may object to it.

Our knowledge of the English language suffices for us to decipher this as a sentence. We know that the symbol 'the' is a definite article denoting shared knowledge and contracting a sense relation with other terms in the determiner system ('the parson' as distinct from 'a parson' or 'this/that parson'). We know that the noun 'parson' denotes a particular religious office, is hyponym to the superordinate 'clergyman' (together with other terms like 'priest', 'rector', 'bishop', and so on). We know how the phrase 'the parson' functions as a constituent, and we identify 'may' as a modal constituent of the verb phrase. With knowledge of this kind, we recognize this expression as a syntactically complete *sentence* and assign it semantic meaning accordingly. But for all this, we do not know what might be meant by the expression in an actual *utterance*, that is to say, when we hear it or read it in a specific context.

Let us imagine somebody coming out with the expression as a remark in the context of a conversation. What kinds of thing might they mean by it? We can *decipher the sentence* by invoking semantic criteria, but how do we *interpret the utterance*?

Reference, force, and effect

Consider the first phrase 'The parson'. The use of the definite article points us in the direction of a particular clergyman assumed to be known about by both speaker and hearer. The noun phrase, then, now takes on a 'pointing' or 'indexical' function, and as such becomes communicatively active as **reference**. But we, as second person parties, have to ratify the reference of course. If we know of no such individual, then the definite article simply directs us into a void, and is indexically invalid ('Parson? Which parson?').

One kind of pragmatic meaning we can assign to an utterance, then, is that of reference. The speaker is talking about something, expressing a **proposition** by using the symbolic conventions of the code to key us into a context of shared knowledge. But the speaker is not just talking about something, but is doing so in order to perform some kind of **illocution** or communicative act. The utterance not only has reference but also **illocutionary force**. So the speaker may intend 'The parson may object to it' as a reason for a decision taken, or as an objection to a particular course of action, or as a warning. 'The parson may object to it.'—'Thanks

for telling me.' These pragmatic possibilities are not signalled in the language itself: they again have to be inferred from the context in which the utterance is made.

One aspect of pragmatic meaning, then, is (propositional) reference, another is (illocutionary) force. There is a third we can identify. In making an utterance, the first person party expresses a certain intended meaning designed not just to be understood as such, but to have some kind of effect on the second person: to frighten, or persuade, or impress, or establish a sense of common purpose or shared concern. ('The parson may object to it.'—'Oh my God!'). This is known as **perlocutionary effect**.

Context and schema

When we talk about propositional reference, illocutionary force, and perlocutionary effect, we are dealing not with the semantic meaning as encoded in the language itself, but the pragmatic meaning which people achieve in **speech acts**.

With speech acts we are again concerned with relationships, but this time not those which are internal to the language itself, but those between aspects of the language and aspects of the external circumstances in which it is used on a particular occasion, its context of occurrence. This context is not just reality in the raw, but those aspects of it which are recognized as significant. Here, we need to invoke again the basic principles of classification which have already been applied. It was pointed out in Chapter 3 that speakers of a language discriminate sounds as phonemically significant by filtering out certain phonetic features. These are not heard as meaningful and so they do not count. In this respect, speakers project their own pattern of reality. The same principle applies to context. When people make an indexical connection, they do so by linking features of the language with familiar features of their world, with what is established in their minds as a normal pattern of reality or **schema**. In other words, context is a schematic construct. It is not 'out there', so to speak, but in the mind. So the achievement of pragmatic meaning is a matter of matching up the linguistic elements of the code with the schematic elements of the context. So, for example, if you were to hear someone make the remark 'Brazil scored just before the final whistle', the likelihood is that the word 'Brazil' would not call to

mind the Amazonian rain forests, coffee, or Copacabana Beach (schematic associations which might be relevant on other contextual occasions), but a football team celebrated for its skill. The football schema thus engaged would lead you to infer what the expressions 'scored' and 'final whistle' referred to among all the possibilities that they *might* refer to in other contexts.

Consider again the comment about the parson. Reference is made to a particular clergyman assumed to be known to both parties. But what is it about him that is relevant here? The 'it' that he might object to could call up the schema associated with his ecclesiastical office: it might refer, for example, to putting a TV in the vestry, replacing the choir with a pop group, using church funds to buy lottery tickets, and so on. What is relevant here is the parson's role as clergyman and custodian of the values of his religion rather than the fact that he is overweight, or near retirement, or unmarried, or plays golf, or rides a bicycle, or smokes a pipe, or whatever. But any one of these *could* be contextually relevant, of course. Everything depends on what 'it' refers to. Reference is achieved when both speaker and hearer engage the same context by converging on what is schematically relevant.

The same thing applies to the achievement of force. The utterance, it was suggested, could be taken as a warning. How might such a force be inferred? Again, the notion of schema comes in. People in a particular community have common assumptions not only about the way the world is organized, but also about the customary ways that social actions like speech acts are performed. It is just these common assumptions that define their cultural identity as members of a social group, small or large. So, in this case, the people in this interaction know that for an utterance to count as the illocutionary act of warning it has to meet certain conditions. To begin with, it obviously has to make reference to some possible future event which would be in some way against the interests of the hearer. But both these conditions apply to the illocutionary act of threat as well. What distinguishes the two is that in the case of a threat, the future event is within the power of the first person to bring about, whereas with a warning it is not.

What then of the parson? If the person whom the utterance is addressed knows that the speaker is on the parson's side, has influence over him, indeed speaks for him, then this will be the

relevant feature about him, the context thereby meets the required condition for the utterance to function as a threat. If, on the other hand, the hearer knows that the speaker does not make common cause with the parson, but sees him as an outside influence over which he, the speaker, has no control, then the utterance will be taken as a warning.

Negotiation of meaning

It may be, of course, that it is unclear whether the context meets one condition or another, whether it is a warning or a threat, and this creates ambiguity. The hearer may eliminate all kinds of circumstantial information as irrelevant to context, but still be left with evidence for more than one possible interpretation. This potential ambiguity applies to all the aspects of pragmatic meaning that we have touched on: reference and effect as well as force. So interpretation commonly involves the parties concerned in the negotiation of meaning, whereby an agreed frame of reference or set of illocutionary conditions is established. One might imagine interactions along the following lines:

A: The parson may object to it.
B: Parson? Which parson?
A: The Reverend Spooner.
B: But he isn't a member. And he doesn't smoke anyway.
A: What's that got to do with the new bicycle shed?
B: I thought you were talking about the smoking ban.

A: The parson may object to it.
B: I don't think we need worry about that.
A: Well perhaps you should. As the chair I must tell you that he will have my support.
B: Yes, and we all know why.
A: That remark is out of order and I must ask you to withdraw it.
B: Don't be such a pompous ass.

A number of other matters arise from these exchanges. Firstly, they are presented here as a written record of an imagined interaction: that is to say, as the text of a supposed discourse. We must assume that many features of such a discourse would remain unrecorded: **paralinguistic** features, for example, like tones of voice, gesture,

facial expression, eye contact, and so on, which might well be contextually relevant and indeed crucial for understanding what is going on. Even if we had recourse to sound tape and video, what would be recorded would be the textual product of the interaction, and not the actual process of the discourse as experienced by the participants.

Secondly, although we began our discussion, as we have in earlier chapters, with simple units of meaning, we are drawn inevitably into a consideration of more complex ones. Although we demonstrated the basic kinds of pragmatic meaning by invoking the speech act as an individual utterance, a kind of pragmatic version of the semantic sentence, it is clear that communication does not take place by the neat sequence of such speech acts. In the first place, they frequently call for negotiation, as we have seen, whereby first person intention and second person interpretation are brought to some satisfactory degree of convergence. Meaning is jointly managed in spoken interaction by turn taking, the sharing of the floor, with different participants assuming the first person speaker role of adjusting the setting for the continuation of the interaction. A major concern of pragmatics is how discourse is managed, what the ground rules for negotiation are, and how (and how far) the different parties cooperate in this joint enterprise. Clearly, when people seek to communicate, they enter into a kind of contract that they will work towards some convergence of intention and interpretation, that is to say, they subscribe to a **cooperative principle**. Otherwise, there would be no way for the semantic potential of language to be given any pragmatic realization at all. There has to be some agreement that what people mean by what they say can be related to what, by established semantic convention, the language itself means. This is not to say that the discourse that people enact will always result in a convergence of *opinion*. Cooperation does not preclude conflict. Indeed, it is only by subscribing to the cooperative principle that people can express disagreement or create conflictual situations.

Relations between utterances

Obviously for any communication to take place, the two parties need to share a common linguistic code (i.e. to speak the same language), but equally they have to be willing, and able, to draw

upon it in accordance with normal communicative conventions. Thus, in our second sample interaction, there is clearly a confrontation developing between A and B. But that very confrontation depends on both parties conforming to the semantic conventions of English as their common code, and also to certain pragmatic conventions which regulate the way the code is used. There are, for example, very general conventions of **cohesion** which establish referential links across the utterances. So it is that each of the interactants recognizes, for example, that 'it' and 'that' refer back to specific things said earlier, that 'you should' and 'we all know why' are reduced or elliptical expressions which are completed by reference to the preceding utterance.

There are also general **turn-taking** conventions which regulate the interaction itself. One of these is the recognition that a pause signals the end of a turn in conversation and an optional shift of speaker role to the second person. Another convention not only constrains the shift of turn but determines what *kind* of turn the second person is to take. So it is, for example, that in asking a question I concede my turn and give you the right to reply. A response is conventionally required. In this respect, question and answer are dependent parts of a single exchange and constitute what is called an **adjacency pair**.

These are very general conventions which regulate the relationship between utterances, but there are more specific ones as well which define how speech acts combine in different modes of communication, or **genres**. The second of the interactions we have been considering, for example, has some of the features of a formal meeting. A convention of this genre is that authority is vested in a chairperson who has the power to control turn-taking and regulate what is said. This accounts for A's statement: 'That remark is out of order and I must ask you to withdraw it.' How, then, do we account for B's reply 'Don't be such a pompous ass'? He may not know the conventions of this particular genre: in other words its formal procedures may not be part of his schematic knowledge. Or it may be that he knows them well enough but chooses to challenge them, seeking to assert a position other than that which A, the chairperson, wants to submit him to.

This illustrates a very general point about pragmatics. It is concerned in part with how language engages the schematic

knowledge people have of what is normal and customary in their particular communities. In this respect pragmatics is the study of how people conform to social conventions. But it is also concerned with the ways such conventions can be circumvented or subverted by individual initiative. Uses of language are, in one respect, necessarily acts of conformity. But they are not only that: they are also acts of identity whereby people assert themselves and manipulate others. Pragmatics is concerned with how people negotiate meaning but also how they negotiate social relations.

And we should note that pragmatics is as much concerned with written as with spoken uses of language. The conventions which come into play for communication and control apply here too. First person writers assume a degree of shared schematic knowledge, produce texts which are cohesive and which conform to the conventions of a particular genre. They count on their readers to cooperate in inferring the values of reference, force, and effect that they intend. To be sure, there can be no immediate reciprocal negotiation of meaning, no joint management of the interaction as there is in conversations. The writer, in sole control, has to make projections about possible readers and anticipate their reactions, thereby enacting a discourse by proxy so to speak, and providing a text as a partial record of it. The readers then have to use the text to activate a discourse of their own, cooperating with the writer as far as they are able or willing to do. In written uses of language, then, the interaction between first and second person parties is displaced and the negotiation of meaning is carried out in two stages. But the meaning is negotiable none the less. It is not inscribed in the language itself and so texts do not signal their own significance. With writing as with speech, pragmatics is concerned with what people *make* of their language.

6
Current issues

Linguistics, like language itself, is dynamic and therefore subject to change. It would lose its validity otherwise, for like all areas of intellectual enquiry, it is continually questioning established ideas and questing after new insights. That is what enquiry means. Its very nature implies a degree of instability. So although there is, in linguistics, a reasonably secure conceptual common ground, which this book has sought to map out, there is, beyond that, a variety of different competing theories, different visions and revisions, disagreements and disputes, about what the scope and purpose of the discipline should be. There are three related issues which are particularly prominent in current debate. One has to do with the very definition of the discipline and takes us back to the question of idealization discussed in Chapter 2. Another issue concerns the nature of linguistic data and has come into prominence with the development of computer programs for the analysis of large corpora of language. A third issue raises the question of accountability and the extent to which linguistic enquiry should be made relevant to the practical problems of everyday life.

The scope of linguistics

As was pointed out in Chapter 2, linguistics has traditionally been based on an idealization which abstracts the formal properties of the language code from the contextual circumstances of actual instances of use, seeking to identify some relatively stable linguistic knowledge (*langue*, or competence) which underlies the vast variety of linguistic behaviour (*parole*, or performance). It was also pointed out that there are two reasons for idealizing to such a

degree of abstraction. One has to do with practical *feasibility*: it is convenient to idealize in this way because the actuality of language behaviour is too elusive to capture by any significant generalization. But the other reason has to do with theoretical *validity*, and it is this which motivates Chomsky's competence–performance distinction. The position here is that the data of actual behaviour are disregarded not because they are elusive but because they are of little real theoretical interest: they do not provide reliable evidence for the essential nature of human language. Over recent years, this formalist definition of the scope of linguistics has been challenged with respect to both feasibility and validity.

As far as feasibility is concerned, it has been demonstrated that the data of behaviour are not so resistant to systematic account as they were made out to be. There are two aspects of behaviour. One is psychological and concerns how linguistic knowledge is organized for access and what the accessing processes might be in both the acquisition and use of language. This has been a subject of enquiry in **psycholinguistics**. The second aspect of behaviour is sociological. This accessing of linguistic knowledge is prompted by some communicative need, some social context which calls for an appropriate use of language. These conditions for appropriateness can be specified, as indeed was demonstrated in part in the discussion of pragmatics in Chapter 5. The account of the relationship between linguistic code and social context is the business of **sociolinguistics**.

Psycholinguistic work on accessing processes and sociolinguistic work on appropriateness conditions have demonstrated that there are aspects of behaviour that can be systematically studied, and that rigorous enquiry does not depend on the high degree of abstraction proposed in formalist linguistics. In other words, psycholinguistics and sociolinguistics have things to say about language which are also within the legitimate scope of the discipline. Such a point of view would be a tolerant and neighbourly one: we stake out different areas of language study, each with its own legitimacy.

But the challenge to the formalist approach in respect to validity is quite different. It is not tolerant and neighbourly at all, but a matter of competing claims for the same territory. It is not just an issue of delimitation but of definition, and proposes a **functionalist**

one in opposition to a formalist one. The argument here is that it diminishes the very study of language to reduce it to abstract forms because to do so is to eliminate from consideration just about everything that is really significant about it and to make it hopelessly remote from people's actual experience. Language, the argument goes, is not essentially a static and well-defined cognitive construct but a mode of communication which is intrinsically dynamic and unstable. Its forms are of significance only so far as we can associate them with their communicative functions. On this account, the only valid linguistics is **functional linguistics**.

But, as was indicated in Chapter 2, there are two senses in which linguistic forms can be said to be associated with functions, and therefore two ways of defining functional linguistics. Firstly, we can consider how the linguistic code has developed in response to the uses to which it is put. In this sense, functional linguistics is the study of how the formal properties of language are *informed* by the functions it serves, how it encodes perceptions of reality, ways of thinking, cultural values, and so on.

Secondly, we can think of the form–function association as a matter not of encoded meaning potential but of its actual realization in communication; and here we are concerned with the way language forms function *pragmatically* in different contexts of use. In this case formalist linguistics is challenged not because it defines the language code too narrowly without regard to the social factors which have formed it, but because it defines language *only* in reference to the code, without regard to how it is put to use in communication. The argument here is that linguistics should extend its scope to account not only for the knowledge of the internalized language of the code, or linguistic competence, but for the knowledge people have of how this is appropriately acted upon, or communicative competence.

These two senses of functional linguistics are frequently confused, and there has sometimes been a tendency to suppose that if you define the code in reference to the communicative functions that have influenced its formation over time, then it follows that you will automatically be accounting for the way in which the code functions in communication here and now. But to do this is to equate the *semantic potential* of the code with actual *pragmatic realizations* of it in communication.

Functional linguistics, in both senses, considers language as an essentially social phenomenon, designed for communication. There is no interest in what makes human language a species-specific endowment, in those universal features of language which might provide evidence of innateness which were described in Chapter 1. The concerns of functional linguistics are closer in this respect to the reality of language as people experience it, and it is therefore often seen as more likely than **formal linguistics** to be applicable to the problems of everyday life. Opponents might argue that this is only achieved at the expense of theoretical rigour. This raises the general question of how far relevance and accountability are valid considerations in linguistic enquiry, and this will be taken up again a little later. It also raises the question of what the source of linguistic data should be, and it is to this matter that we now turn.

The data of linguistics

There are, broadly speaking, three sources of linguistic data we can draw upon to infer facts about language. We can, to begin with, use *introspection*, appealing to our own intuitive competence as the data source. This is a tradition in linguistics of long standing, and essentially makes operational Saussure's concept of *langue* as common knowledge, imprinted in the mind like a book of which all members of the community have identical copies. So if linguists want data, as representative members of a language community they have only to consult the copy in their head. Most grammars and dictionaries until recent times have been based on this assumption that linguistic description can be drawn from the linguist's introspection. And it is not only linguistic competence which is accessible to introspection, but communicative competence as well, so the argument is that the conventions that define appropriate language use can also be drawn from the same intuitive source.

If, however, there is some reason to doubt the representative nature of such intuitive sampling, there is a second way of getting at data, namely by *elicitation*. In this case, you use other members of the community as informants, drawing on *their* intuitions. And again, this might be directed at obtaining the data of the code or

its communicative use. Thus, you might ask informants whether a particular combination of linguistic elements are grammatically possible in their language, or what would be an appropriate expression given a particular context.

Introspection and elicitation can be used to establish both the formal properties of a language and how they typically function in use. But in both cases the data is abstract knowledge, and not actual behaviour. They reveal what people know about what they do but not what they actually do. If you want data of that kind, the data of performance rather than competence, you need to turn to *observation*.

The development of computer technology over recent years has made observation possible on a vast scale. Programs have been devised within **corpus linguistics** to collect and analyse large corpora of actually occurring language, both written and spoken, and this analysis reveals facts about the frequency and co-occurrence of lexical and grammatical items which are not intuitively accessible by introspection or elicitation.

It would seem on the face of it that this is a much more reliable source of data. It is surely better to find out what people actually do than depend on intuitions which are often uncertain and contradictory. Claims have indeed been made that these large-scale observations reveal patterns of attested usage which call for a complete revision of the existing categories of linguistic description, which are generally based on intuition and elicitation. Corpus linguistics, in dealing with actual behaviour, clearly has an affinity with functional linguistics in that it too claims to get closer to the facts of 'real' language.

There is no doubt that corpus analysis can reveal facts of usage, the data of actual linguistic performance, which throw doubt on the validity of any model of language based on the idea of a stable and well-defined system. The elaborate picture it presents is very different from the abstract painting proposed by the formal linguist. If language use is indeed a rule-governed activity, as is often said, the rules are not easy to discern in the detail. And it is also true that this detail is not accessible to introspection or elicitation. Even a limited corpus analysis can show patterns of occurrence of which language users, the very producers of the data, are unaware. Corpus linguistics transcends intuitive knowledge and

in this respect can be seen as a valuable, and valid, corrective to unfounded abstraction: a case of description influencing theory for once, rather than the other way round.

But the claims of corpus linguistics can be questioned too. The facts of usage revealed by computer analysis, for example, carry no guarantee of absolute truth. The intuitions that people have about their language have their own validity as data. These conceptual constructs are also real, but the reality is of a different order.

One example of this is the way lexical knowledge (in some areas of vocabulary at least) seems to be organized semantically in terms of **prototypes**, and these cannot be observed, but only elicited. Thus, when a group of English-speaking informants were asked to give the first example that came to mind of a more inclusive category of things they showed a striking unanimity. The word 'bird' elicited 'robin' (rather than, say, 'chaffinch' or 'wren') and the word 'vegetable' elicited 'pea' (rather than, say, 'parnsip' or 'potato'). For these informants, then, a robin is the prototypical bird, a pea the prototypical vegetable. But this conceptual preference does not correspond with how frequently these words actually occur in a corpus. The same point can be made about grammatical structures. If English-speaking informants are asked to provide examples of a sentence, they are likely to come up with simple subject–verb–object (SVO) constructions ('The man opened the door'; 'John kissed Mary'). These, we might say, are prototypical English sentences. But they are unlikely to figure very frequently in a corpus of actual usage. Since people do not use simple sentences like this very often, they do not have much reality as observed data, but they may have a significant psychological reality nevertheless. They may be evidence of competence which is not reflected in the facts of performance.

Prototypes thus elicited do not, of course, invalidate the observed data of corpus linguistics. They provide a different kind of data which are evidence of competence which is not directly projected into performance. Intuitive, elicited, and observed data all have their own validity, but this validity depends on what kind of evidence you are looking for, on what aspects of language knowledge or behaviour you are seeking to explain. If you are looking for evidence of the internal relationship between lan-

guage and the mind, you are more likely to favour intuition and elicitation. If you are looking for evidence of how language sets up external links with society, then you are more likely to look to the observed data of actual occurrence. The validity of different kinds of linguistic data is not absolute but relative: one kind is no more 'real' than another. It depends on what you claim the data are evidence of, and what you are trying to explain.

The relevance of linguistics

From questions of validity we turn now to questions of utility. What is linguistics *for*? What good is it to anybody? What practical uses can it be put to? One response to such questions is, of course, to deny the presupposition that it needs any practical justification at all. Like other disciplines, linguistics is an intellectual enquiry, a quest for explanation, and that is sufficient justification in itself. Understanding does not have to be accountable to practical utility, particularly when it concerns the nature of language, which, as was indicated in Chapter 1, is so essential and distinctive a feature of the human species.

Whether or not linguistics *should* be accountable, it has been turned to practical account. Indeed, one important impetus for the development of linguistics in the first part of this century was the dedicated work done in translating the Bible into languages hitherto unwritten and undescribed. This practical task implied a prior exercise in **descriptive linguistics**, since it involved the analysis of the languages (through elicitation and observation) into which the scriptures were to be rendered. And this necessarily called for a continual reconsideration of established linguistic categories to ensure that they were relevant to languages other than those, like English, upon which they were originally based. The practical tasks of description and translation inevitably raised issues of wider theoretical import.

They raise other issues as well about the relationship between theory and practice and the role of the linguist, issues which are of current relevance in other areas of enquiry, and which bear upon the relationship between descriptive and **applied linguistics**.

The process of translation involves the interpretation of a text encoded in one language and the rendering of it into another text

which, though necessarily different in form, is, as far as possible, equivalent in meaning. In so far as it raises questions about the differences between language codes it can be seen as an exercise in **contrastive analysis**. In so far as it raises questions about the meaning of particular texts, particular communicative uses of the codes, it can be seen as an exercise in **discourse analysis**. Both of these areas of enquiry have laid claim to practical relevance and so to be the business of applied linguistics.

With regard to contrastive analysis, one obvious area of application is language teaching. After all, second language learning, like translation, has to do with working out relationships between one language and another: the first language (L1) you know and the second language (L2) you do not. It seems self-evident that the points of difference between the two codes will constitute areas of difficulty for learners and that a contrastive analysis will therefore be of service in the design of a teaching programme.

It turns out, however, that the findings of such analysis cannot be directly applied in this way. Although learners do undoubtedly refer the second language they are learning (L2) to their own mother tongue (L1), in effect using translation as a strategy for learning, they do not do so in any regular or predictable manner. Linguistic *difference* is not a reliable measure of learning *difficulty*. The data of actual learner performance, as established by **error analysis**, call for an alternative theoretical explanation.

One possibility is that learners conform to a pre-programmed cognitive agenda and so acquire features of language in a particular order of acquisition. In this way they proceed through different interim stages of an **interlanguage** which is unique to the acquisition process itself. Enquiry into this possibility in **Second Language Acquisition (SLA)** research has been extensive.

There is another possibility. It might be that the categories of description typically used in contrastive analysis are not sufficiently sensitive to record certain aspects of learner language. Learners may be influenced by features of their L1 experience other than the most obvious forms of the code. Contrastive analysis has been mainly concerned with syntactic structure, but as we have seen in Chapters 4 and 5, this is only one aspect of language, and one which, furthermore, inter-relates with others in complex ways. So it may be that the learners' difficulties do correspond to

differences between their L1 and L2, but that we need a more sophisticated theory to discern what the differences are, a theory which takes a more comprehensive view of the nature of language by taking discourse into account.

Discourse analysis is potentially relevant to the problems of language pedagogy in two other ways. Firstly, it can provide a means of describing the eventual goal of learning, the ability to communicate, and so to cope with the conventions of use associated with certain discourses, written or spoken. Secondly, it can provide the means of describing the contexts which are set up in classrooms to induce the process of learning. In this case it can provide a basis for classroom research.

But the relevance of discourse analysis is not confined to language teaching. It can be used to investigate how language is used to sustain social institutions and manipulate opinion; how it is used in the expression of ideology and the exercise of power. Such investigations in **critical discourse analysis** seek to raise awareness of the social significance and the political implications of language use. Discourse analysis can also be directed to developing awareness of the significance of linguistic features in the interpretation of literary texts, the particular concern of **stylistics**.

In these and other cases, descriptive linguistics becomes applied linguistics to the extent that the descriptions can be shown to be relevant to an understanding of practical concerns associated with language use and learning. These concerns may take the form of quite specific problems: how to design a literacy programme, for example, or how to interpret linguistic evidence in a court of law (the concern of the growing field of **forensic linguistics**).

But other concerns for relevance are more general and more broadly educational. We began this book by noting how thoroughly language pervades our reality, how central it is to our lives as individuals and social beings. To remain unaware of it what it is and how it works is to run the risk of being deprived or exploited. Control of language is, to a considerable degree, control of power. Language is too important a human resource for its understanding to be kept confined to linguists. Language is so implicated in human life that we need to be as fully aware of it as possible, for otherwise we remain in ignorance of what constitutes our essential humanity.