CHAPTER

1 Meaning in the empirical study of language

CHAPTER PREVIEW

In this chapter we will introduce some important concepts for the study of semantics. In 1.1 we place the notion of linguistic meaning in the wider context of human communication and behaviour. Section 1.2 then examines some of the vocabulary that English and other languages use for ordinary talk about meaning in language and related phenomena. A consideration of how this everyday non-technical vocabulary varies cross-linguistically can show some of the important different aspects of linguistic meaning. In section 1.3 the **semiotic triangle** of mind, world and language is discussed, followed in 1.4 by an introduction to five fundamental concepts:

- lexemes;
- sense and reference;
- denotation and connotation;
- compositionality; and
- levels of meaning.

Next (1.5), we introduce the concepts of **object language** and **metalanguage**, and distinguish a number of different possible relations between the language *in which* meanings are described (the 'metalanguage') and the language *whose* meanings are described (the 'object language'). We will then consider three different identifications of meaning: meanings as objects in the world (referents: 1.6.1), as objects in the mind (concepts: 1.6.2), and as brain states (1.6.3). An alternative identification is the notion of meanings as uses, discussed in 1.6.4. To end the chapter, we consider a view of meaning on which meanings are unobservable, hypothetical constructs posited to explain facts about language use (1.7).

1.0 What is semantics?

Any attempt to understand the nature of language must try to describe and explain the ways in which linguistic expressions have meaning. This book introduces some of the aspects of meaning studied in linguistic semantics, the branch of linguistics which, along with pragmatics, has responsibility for this task. Semantics is one of the richest and most fascinating parts of linguistics. Among the kinds of questions semanticists ask are the following:

- What are meanings — definitions? ideas in our heads? sets of objects in the world?
- Can all meanings be precisely defined?
- What explains relations between meanings, like synonymy, antonymy (oppositeness), and so on?
- How do the meanings of words combine to create the meanings of sentences?
- What is the difference between literal and non-literal meaning?
- How do meanings relate to the minds of language users, and to the things words refer to?
- What is the connection between what a word means, and the contexts in which it is used?
- How do the meanings of words interact with syntactic rules and principles?
- Do all languages express the same meanings?
- How do meanings change?

Clearly, semantics is a vast subject, and in this book we will only be able to introduce the most important parts of it. 'Meaning', however, is a very vague term. In ordinary English, the word 'meaning' is used to refer to such different things as the *idea* or *intention* lying behind a piece of language, as in (1), the *thing referred to* by a piece of language (2), and the translations of words between languages (3).

(1) *'I don't quite understand what you're getting at by saying "meat is murder": do you mean that everyone should be a vegetarian?'*

(2) *'I meant the second street on the left, not the first one.'*

(3) 'Seiketsu *means "clean" in Japanese.*'

As we will see, an important initial task of linguistic semantics is to distinguish between these different types of meaning, and to make it clear exactly what place each of them has within a principled theory of language (see Sections 1.4 and 1.6).

Each of the chapters of this book introduces some essential concepts for understanding the ways in which meaning can be analysed in linguistics. This first chapter is an introduction to the issues and concepts studied in linguistic semantics. In Chapter 2 we consider the relation between

meanings and definitions. When we think about word meanings, definitions in dictionaries quickly come to mind: we know that, if uncertain about a word's meaning, we can look it up in a dictionary. This means that it is important to be clear about the similarities and differences between the aspects of meaning that interest linguists, on the one hand, and lexicographers (dictionary-writers) on the other. In Chapters 3 and 4 we discuss the relation between word meaning and word use: how do we distinguish between what a word actually means, and the way in which it happens to be used on a given occasion? Chapter 5 looks at attempts to analyse the meanings of words into sets of basic components, and discusses the problem of determining just how many meanings a given word has. In Chapter 6 we introduce some concepts from formal logic which have been fruitfully applied to the analysis of natural language meanings, and in Chapters 7 and 8 we look at the ways research inspired by psychology has been used to illuminate linguistic semantic questions, and how the results of this research can be modelled on computers. Chapter 9 explores the semantics of the parts of speech and of tense and aspect. Chapter 10 discusses the relationship between semantics and syntax, a subject which raises many important questions. Chapter 11 emphasizes a somewhat different aspect of meaning, its changeability. Meaning is always changing, both synchronically (i.e. between different speakers at the same time) and diachronically (over time). No comprehensive study of meaning can neglect this variation and change.

QUESTION How closely does the subject matter of semantics seem to correspond with what you would have thought are the main questions to ask about meaning in language?

1.1 Meaning, communication and significance

Informally, it is easy to agree that meaning is the heart of language. Meaning, we might say, is what language is *for*: to have a language without meaning would be like having lungs without air. Only when sequences of sounds or letters have (or are judged capable of having) a meaning do they qualify as language: infants' babbling and bird song, for example, use the same medium as human language – sound – but since they do not, and cannot, express meaning (except, perhaps, to the infants or the birds) we do not consider them as examples of language in the full sense of the word. Meaning is also central to the *experience* of using language, as anyone knows who has ever listened to people talking in an unknown language. Not only does such a language fail to express any meaning; it is also often hard to catch hold of individual words: without knowing the meaning of an utterance, it is hard to identify the separate words which constitute it.

Without a capacity to express meaning, then, language loses one of its essential aspects. We practically always speak or write in order to express a meaning of one kind or another. This is most obviously true for pieces

of language which convey information: if someone suddenly says (4), then a meaning has been conveyed, and you are in possession of some information – whether true or false – which you may not have previously known.

(4) *Engels was two and a half years younger than Marx.*

But not only sentences have meanings. Even the shortest, most everyday words, which we would not normally consider as containing information, like *the*, *not, of,* or even *ouch!*, contribute something specific to the meanings of utterances in which they occur and can thus be legitimately considered as having meanings in their own right. (For some scholars, the study of the meanings of words like these belongs as much to pragmatics and syntax as it does to semantics; we will discuss the difference between semantics and pragmatics in 1.4.4.)

Question Two apparent exceptions to the meaningfulness of language are T-shirts worn in Japan and elsewhere with 'nonsensical' English sentences on them, and people speaking in tongues at certain religious meetings. Are there other examples of this kind? Are instances of language use like this really non-meaningful? If so, what are some possible implications for semantics? If not, why not?

Although the study of meaning is extremely ancient, the name *semantics* was only coined in the late nineteenth century by the French linguist Michel Bréal. Like many other names of branches of linguistics, the word *semantics* reflects the origins of the Western tradition of linguistic analysis in the writings of Greek thinkers from the fifth century BC onwards. *Semantics* comes from the ancient Greek word *semantikos*, an adjective meaning 'relating to signs', based on the noun *sēmeion* 'sign'. In Ancient Greek, one of the original uses of *sēmeion* was as a medical term for the symptoms that were the *signs* of underlying diseases. This derivation highlights the close relation between the study of linguistic signs – words, phrases, sentences and utterances – and the study of signs in general: both artificial, conventional signs like road signs, clock faces, the symbols used in computer programs, or the 'signals' communicated by different choices of clothes; and natural signs like symptoms of disease, the level of the sun in the sky (a sign of the time of day) or tracks on the ground (the sign that an animal has passed). The study of signs in general is known as **semiotics** or **semiology** (both Greek words also deriving from *sēmeion*). In the twentieth century, the general study of signs became particularly important and the new discipline of semiotics was created, especially as the result of the work of the American philosopher Charles Sanders Peirce (pronounced 'purse'; 1839–1914) and of Bréal's student, the Swiss linguist Ferdinand de Saussure (1857–1913), often considered as the founder of modern linguistics.

The meanings we can express through language are infinitely more numerous, detailed and precise than those expressible through other semiotic media. Yet the type of meaning found in language can be seen as a subset of two broader categories of meaningfulness: the significance of

human behaviour in general, and the meaningfulness of communication specifically. There are many meaningful ways of behaving which do *not* involve language. These are not limited to those types of behaviour involving structured sets of conventional, accepted symbols like the left-right indicator lights on cars, the use of flags at sea to convey various specific messages, or the many types of symbol involving body parts (bowing, waving, nodding and shaking the head, the thumbs up/thumbs down signals, the hand signs used in baseball, etc.). Many types of intentional human behaviour can be seen as having a significance, or a meaning, in the (broad) sense of the word, since they both express, and allow observers to draw conclusions about, the nature and intentions of the participants. Someone who has just got up from their seat on the bus is probably intending to get off. Someone who suddenly stops walking down the street to search frantically through their pockets may just have realized that they have forgotten their keys. Unlike the use of language, these types of behaviour do not involve any structured set of symbols or, necessarily, any **communicative intention** and are therefore non-semiotic. The person getting up from their seat is not wishing to communicate anything to anyone, and is not making use of any structured communicative symbols: they simply want to get off. The use of fully articulated language, which does involve a communicative intention, is thus only the fullest and most explicit way in which we derive information about our environment: as a result, the meaningfulness of language can be seen as a subset of the meaningfulness of human behaviour.

QUESTION We have just given a number of examples of conventional symbols. What are some others?

Even when an intention to communicate does exist, however, the use of language is only one of a number of ways in which the intention can be fulfilled. Take the example of someone at the dinner table suddenly choking on some food. They start to gasp, they go red in the face, their eyes water, and all they can do is make a muffled, indistinct cry. To the other people at the table, this communicates something: they realize that there is something wrong and that help is needed. As a result, they could quickly help the sufferer by giving them a glass of water or a slap on the back. This, then, is an example of some information being made known without the help of language: the person choking has just cried out, perhaps involuntarily, and this is enough to attract the attention of others, to tell them something about the current state of that person, and to stimulate them to bring the required help. Now imagine that the person choking, instead of simply crying out, articulates three quick syllables consisting simply of three choking-sounding vowels, with the middle syllable louder than the others: '*-*-*'. In this case, the other people at the table might conclude that the three cries were substitutes for the three syllables of the sentence 'I'm CHOking!', and would act on the basis of this (correct) assumption. Here, even though the speaker can only manage to articulate the syllable pattern of the intended phrase, communication

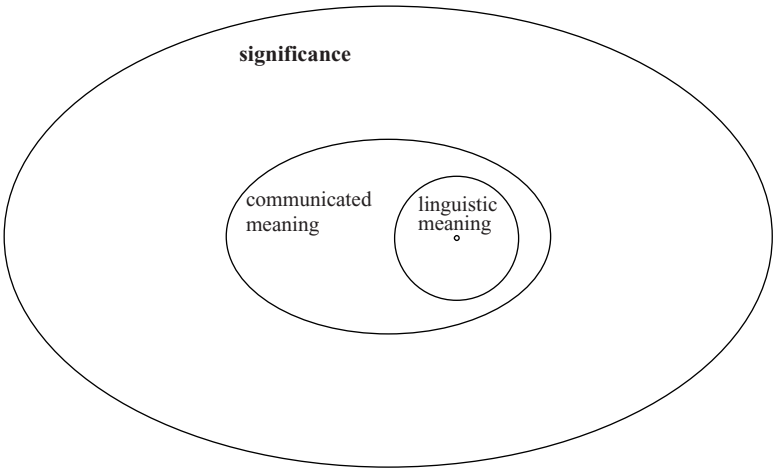

FIGURE 1.1
Significance, communicated meaning and linguistic meaning.

successfully takes place. Of course, if they had enough breath left, they could simply cry out 'I'm choking', and there would be no ambiguity. These cases show that a fully articulated sentence is not always necessary to communicate an intended meaning: the same meaning can be suggested in a variety of different ways, all of which rely on implicit conventions. The sentence expresses the intended meaning more precisely and unambiguously than the others: both the single cry and its three syllable variant are open to many interpretations, and are therefore much less reliable than the fully explicit sentence. But we can nevertheless remove the language from a communicative situation and retain much of the meaning. Situations are inherently meaningful. Meaning, we might say, is already there in the world: all we have to do is draw attention to it, and language is the most specific and unambiguous way of doing so. The different types of meaningfulness we have been discussing so far could be diagrammed as in Figure 1.1.

1.2 Talking about meaning in English and other languages

Semantics, then, is the study of meaning. But what actually *is* meaning? In Section 1.6 we will discuss some specific answers to this question. For the moment, we will make a start by looking at what place the notion of meaning has in our ordinary talk about language. The way we use the concept of meaning in ordinary language is important because it provides us with a *pretheoretical* starting point for theoretical semantic analysis, and gives us the initial vocabulary with which we can begin to identify and describe the phenomena which strike us. Informal talk about what pieces of language mean is a very common part of everyday life: we explain new words, give paraphrases of what people mean by a certain phrase or expression, sometimes translate words from one language to another in order to show their meaning. But even though we *use* the

notion of meaning naturally and unproblematically, it is quite another thing to develop an explicit, rigorous *explanation* of it. In just the same way, it is one thing to talk about the movements of celestial bodies like the moon and stars – we do so, informally, all the time – but a different one entirely to have a scientific understanding of them. And since meanings cannot be seen, there is the initial question of how to pin down exactly what we are and are not supposed to be investigating. It will help us to accomplish this task if we examine the everyday vocabulary used to talk about meaning in English and other languages. This vocabulary varies considerably cross-linguistically; examining it can show some of the important different aspects of linguistic meaning, and can allow us to see how different languages impose different starting distinctions on what we, in English, call 'meaning'.

1.2.1 'Meaning' in English

English uses the verb *to mean* to refer to a relationship involving at least one of three different types of thing: language, the world (including people, objects, and everything outside of ourselves) and our own minds or intentions. Here are five typical examples of *mean* in English which exemplify some of these relationships:

(5) *When I said 'Dublin has lots of attractions' I meant Dublin, Ireland, not Dublin, Virginia.*

(6) *In Sydney, 'the bridge' means the Harbour Bridge.*

(7) *'Stout' means 'short and fat'.*

(8) *By turning off the music I didn't mean that you should go.*

(9) *Trees mean water.*

Sentence (5) distinguishes two possible places that the speaker could have been referring to by the name 'Dublin', and specifies that only one of them was intended. This, then, is a three-way relation between a piece of language, a mind and the world: the world is represented by the two places called Dublin, language by the sentence 'Dublin has lots of attractions', and mind by the speaker's *intention* to refer to Dublin, Ireland. The second sentence is a relation between language and world, without any specific reference to people's intentions. It says that the expression 'the bridge' refers to one particular structure – the Sydney Harbour Bridge – rather than any of the other bridges in Sydney. Even though it is obviously only through the action of speakers' minds that *bridge* has this reference, there is no explicit mention of speakers' minds in (6). In (7), there is no explicit reference to either people's minds or to the world: the sentence reports an equivalence between two linguistic items, the word 'stout', according to (7), is simply equivalent in some way to the words 'short and fat'. Sentence (8) refers to a mind–world relation: it is thus like sentence (5), except that there is no language: the speaker denies that the action of turning the music off was the result of any *intention* for the guests to leave.

Sentence (9) names a world–world relationship: the presence of one type of object in the world (trees) reveals the presence of another (water).

The fact that the same verb is used in English for these non-linguistic situations as well as the linguistic ones is noteworthy if we consider the discussion in 1.1. Thus, while sentences (5)–(7) refer to linguistic meaning, sentence (8) refers to communicated meaning, and sentence (9) refers to what we have called significance. In sentence (8) (spoken, say, at a party where it has got late and there are only a few guests left), the act of turning off the music could be interpreted as a sign of the end of the party: sentence (8) is a way of saying that the speaker did not intend this. And to say that 'Trees mean water' is to say that the presence of trees allows us to conclude that there must be water nearby (compare the examples of significance in the previous section). This is a conclusion we reach simply by virtue of what we know about trees and water, and without there being any communication as such.

In ordinary English, then, we use the same verb to refer both to the meanings expressed by language and to those which are communicated non-linguistically, as well as to those which emerge, without any communication, as a result of the inherent significance of the world and human behaviour. In a number of these situations, the idea of the intention of the communicator seems to be an important part of what is being talked about through the use of the verb *mean*. But meaning is not the only way in which situations like those in (5)–(6) can be described in English: a number of other possible modes of description are also available. To see this, let's narrow the discussion down to one particular example of language – a piece which many people would think of as, simply, a mistake. Consider the following situation: Judy and Alastair are having a dinner party, and Alastair has gone out to buy a few extra plates and cups for the guests. Coming home, he says:

(10) *I've got some more cutlery for the party.*

For most speakers of English, this would count as a mistake, since 'cutlery' refers not to cups and plates, but to knives, forks and spoons. But the fact that this is a mistake in no way diminishes the need for a principled, linguistic account of it: like other branches of linguistics, semantics describes language as it is actually used and the use of a mistake as our example here will allow the relevant issues to emerge particularly clearly.

How then can we describe what is happening in (10)? In context, we can imagine three replies which Judy might make, each of which considers Alastair's 'mistake' from a different point of view:

(11) a. Judy: *Cutlery?! We've got lots of cutlery! You mean you got more* crockery*!*
Alastair: *Oh yeah, crockery.*

b. Judy: *Cutlery?! Why did you say* cutlery *instead of* crockery*?*
Alastair: *Oh yeah, crockery.*

c. Judy: *Cutlery?! You did not! You got more crockery!*
Alastair: *Oh yeah, crockery.*

In (11a) Judy uses the category of meaning to describe Alastair's language, and says that Alastair did not actually *mean* 'cutlery': what he meant was 'crockery'. In (11b) she talks about what Alastair 'says'. Here, she could be described as talking not about language meaning, but language *use*: she notes that Alastair has used the term *cutlery* when the term *crockery* would be expected. In (11c), Judy simply denies what Alastair has said. In so doing, she can be described as applying the categories of truth and falsity to Alastair's utterance: according to her, it is simply not true that Alastair bought cutlery, a fact which Alastair then admits.

Ordinary English, then, makes available at least three different ways of talking about language: meaning, use and truth. Each of these three categories of ordinary language description highlights a particular aspect of the occurrence. Description in terms of truth places the emphasis on the objective facts of the situation by concentrating on the relation between language and reality: does the language used correspond to the actual state of affairs? Description in terms of use makes no explicit reference to the facts, but limits itself to a consideration of equivalences between the piece of language in question and an assumed norm: Alastair said *cutlery* when, in the same circumstances, most people would have said *crockery*. Lastly, description in terms of meaning places the emphasis on the speaker's intentions: for Judy to say that Alastair meant crockery is, in this context, the equivalent of saying that he *intended* to say *crockery*, and to note a discrepancy between this assumed intention and the actual words used.

As we will see in Section 1.6, each of these ordinary language modes of description has its own developed, theoretical analogue.

1.2.2 'Meaning' in Warlpiri

In English, then, the one verb 'mean' is used to describe reference, linguistic meaning, intention, and general significance. Given the frequency with which, in English, we use this verb to talk about the relations between language, intention and the world, it may be surprising to discover that there are languages which do not make use of any similar notion in order to talk about situations like those in (5)–(6) above. One such language is Warlpiri, a Pama-Nyungan language spoken in central Australia. In a sense, Warlpiri has no equivalent for the verb *mean*, and the links between reference, linguistic equivalence, intention, and general significance are quite differently constituted.

In Warlpiri, the most common way of asking about the 'meaning' of a word does not involve any verb. For example, to ask about the meaning of the word *karnta* ('woman'), one would simply say (12):

(12) *Nyiya* *karnta-ju?*
what *karnta*-TOPIC
'What is a *karnta*?'/'What does "*karnta*" mean?'

This could be translated as either 'what does *karnta* mean?' or as 'what is a *karnta*?'. And when the meaning of a word is explained or defined, once

again no separate verb meaning 'mean' is involved. In the following example, for instance, the speaker is explaining the meaning of the word *ngalyarra*:

(13) *Ngalyarra* *ngula-ju* *yanjilypiri* *panu.*
Ngalyarra that-TOPIC stars many
'Ngalyarra – that is many stars'/'*Ngalyarra* means "many stars".'
(WlpD: *ngalyarra*)

The absence of the specific verb 'mean' is characteristic of a wider set of contexts in Warlpiri; there is also very often no separate verb that would be the equivalent of 'is' in English, as the following examples show:

(14) *Ngamirliri,* *ngula-ji* *kirrirdipardu.*
curlew that-TOPIC tall
'The curlew is tall.' (WlpD: *ngamirliri*)

(15) *Jajirdi* *kuyu* *wita.*
native cat animal small
'The native cat is a small animal.' (WlpD: *jajirdi*)

The result of this is that Warlpiri makes less of a distinction than English between what a *word* means, and what its *referent* actually *is*. To say what a word means is simply to describe the object or situation it refers to. Language–world relations are described in the same way as world–world ones.

Warlpiri does, however, have a way of explicitly mentioning the language-user, as can be seen in the following example:

(16) *Mirni-nya* *karnalu* *wurnturu* *ngarri-rni.* *Kala* *mirnimpa,*
mirni-FOCUS 1PL.SUBJ far call-NONPAST but *mirnimpa*
ngula-ju *kutu-pardu* *karnalu* *ngarri-rni.*
that-TOPIC close-rather 1PL.SUBJ call-NONPAST
'We use *mirni* to mean far, whereas by *mirnimpa* we mean rather close.' (WlpD: *mirnimpa*)

But the verb used here, *ngarri-rni*, which simply means 'call', does not make any reference to the speaker's intentions, an important component of the notion of 'meaning' in English. The literal meaning of (16) is something like 'we call far things *mirni*, whereas we call close things *mirnimpa*.' This is simply a fact about language use: *ngarrirni* 'call' makes no reference to any intention of the speaker, and the verb *manngi-nyanyi* 'think, intend', is not typically used to refer to the meaning of words.

1.2.3 'Meaning' in French

Whereas, in Warlpiri, the meanings of words are not discussed in the same terms as the intentions of speakers, in French there is a close link between these two domains. The most common way of expressing 'mean'

in French is the expression 'vouloir dire', which literally means 'to want to say.' To ask 'what do you mean?' in French is to ask 'what do you want to say?' Talking about meaning in French, then, inherently involves talking about volition ('wanting'), as in the following expressions:

(17) *Qu'est-ce* *que* *tu* *veux* *dire* *par* *cela?*
what is it that you want to say by that?
'What do you mean by that?'

(18) *Que* *veut* *dire* *cette* *phrase* *latine?*
what wants to say this phrase latin
'What does this Latin phrase mean?'

(19) *Que* *veut* *dire* *ce* *vacarme,* *cette* *agitation?*
what wants to say this clamour this agitation
'What does this clamour and agitation mean?'

(20) *Le baromètre a* *baissé;* *cela* *veut* *dire* *qu'* *il*
the barometer has gone down that wants to say that it
va *pleuvoir*.
is going to rain
'The barometer has gone down; that means it's going to rain.'

As (19) and (20) show, this is even the case when talking of what words, phrases and non-linguistic things mean: as in English, the same expression is used to refer both to the meaning of language, and the meaning of non-linguistic occurrences. *Vouloir dire* is not, of course, the *only* word available in French for the expression of ideas about meaning; the verb *signifier* (from the Latin *signum* 'sign' and *facere* 'to make') has a similar sense. Another contrast between French and English is that unlike in English, the French words that express the noun 'meaning' and the verb 'to mean' are not related. In French the noun 'meaning' is translated by the word *sens*, from which English gets the word 'sense', and which has a similar range of meanings: as well as referring to linguistic meaning, *sens* refers to the perceptual senses (sight, hearing, etc.), to a direct and intuitive grasp of something (e.g. a 'sense' of rhythm), as well as having the meaning expressed in English by saying that something 'makes sense'. Just like *vouloir dire*, then, *sens* classes linguistic meaning together with certain inner, subjective processes of human consciousness; not, however, as in the case of *vouloir dire*, volitional ones, but ones connected with the faculties of perception and judgement.

1.2.4 'Meaning' in Chinese

In Mandarin Chinese, there is no single word with the same range of meanings as English *mean* or *meaning*. The verb *zhi*, whose core meaning is 'point', can express all of the relations between mind, language and world discussed in the previous sections, except the world–world relation. Thus, we find *zhi* used for the mind–language–world relation, as in (21):

(21) *Dang* *wo* *shuo* *'Coles',* *wo* *shi* *zhi* *Central* *de*
when I say 'Coles' I BE point Central POSS
'Coles', *bu* *shi* *TownHall* *de* *'Coles'.*
'Coles' not BE TownHall POSS 'Coles'
'When I say "Coles", I mean the "Coles" in Central but not the "Coles" in Town Hall.'

As well, it can be used for the language–world relation:

(22) *Zao-can* *shi* *zhi* *zao-shang* *chi* *de* *yi* *can.*
breakfast BE point morning eat POSS one meal
'"Breakfast" means the meal you have in the morning.'

Zhi may also be used to specify a word's translation:

(23) *'Linguistics'* *shi* *zhi* *yu-yan-xue.*
'Linguistics' BE point yu-yan-xue
'"Linguistics" means *yu-yan-xue*.'

However, when a monolingual definition is given, the noun *yi-si* 'meaning' is typically used:

(24) *Miao-tiao* *de* *yi-si* *shi* *shou* *ji* *xian-xi*
'Miao-tiao' POSS meaning BE thin and delicate
'"*Miao-tiao*" means thin and delicate.'

Yi-si is also used in a way that parallels the English use of *meaning* to express the language–mind relation:

(25) *Wo* *ming-bai* *ne* *de* *yi-si.*
I understand you POSS meaning
'I understand what you mean.'

A native speaker explains *yi-si* here in the following way: 'the speaker is conveying the message that he can reveal what's in the hearer's mind and the intention behind it. It is actually similar to saying "I understand what you are thinking about"' (W. Chor, p.c.). But *yi-si* cannot be used for the world–world relation:

(26) **Jin-qian de* *ji-si* *shi quan-li.*
money-POSS meaning BE power
'Money means power.'

To express this, *deng-yu* 'equal' may be used:

(27) *Jin-qian* *deng-yu* *quan-li.*
money equal power
'Money means power.'

We thus find that, taken together, the translations of *mean/meaning* in Mandarin have a similar range of senses to their English equivalents, except that Mandarin has no equivalent to *money means power* or *clouds mean rain*. However, the fact that the verb meaning 'point' is the basic way of expressing the verbal notion brings in a connection between meaning and gesture which is not familiar from English.

1.3 The semiotic triangle: language, mind, world and meaning

We have seen in the previous section that a number of languages, including French and English, make an important connection in their standard vocabularies between language and the world of inner conscious processes like volition, perception and intention. Other languages, by contrast, like Warlpiri, seem to bypass this connection by talking about the meaning of language in the same terms used to talk about the identity of things in the world. All of these relations are important. To describe meaning fully, we seem to have to make reference to three principal terms: language, the world, and the human mind. Following Ogden and Richards (1949: 10), these three aspects of the meaning phenomenon are often symbolized as the 'semiotic triangle', as shown in Figure 1.2 below.

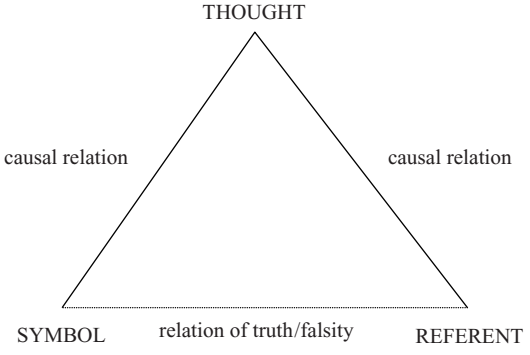

FIGURE 1.2
The semiotic triangle.

At the top of the triangle is what Ogden and Richards called 'thought'. This reflects the fact that language comes from human beings, and is therefore ultimately a product of processes in the mind or brain. But 'thought' can be a misleading label for these processes, for two reasons. First, these mental processes need not be conscious. Even though we sometimes do consciously think about what we are going to say, our speech is more often spontaneous, emerging without our being aware of any preliminary stage of mental preparation. Since it is the brain that produces language, we know that some such preliminary stage must have taken place, but since this stage is so often unconscious, the label 'thought' is not the most appropriate (see Chapter 11 for more discussion).

The second reason that 'thought' is an unfortunate label for the mental processes at the origin of speech is that it excludes the non-rational, emotional side of our inner life. The processes leading to speech should not be limited to what we would class simply as 'thinking', but extend to include our emotions and volition as well. This is most obviously true with exclamations: exclamations of pain, surprise or happiness often do not reflect anything we would describe as a 'thought', but rather reflect a particular feeling. The same is true for many other types of words, like diminutives, which may correspond to feelings of (roughly) affection; and imperatives, which may be accompanied by feelings of control, superiority, pride, etc. Evaluative words more generally, expressing the speaker's emotional attitude, often force us to recognize a strong emotional component. Thus, 'marvellous', 'wonderful', 'fantastic' and 'good'; and 'appalling', 'terrible', 'frightful' and 'bad' and their synonyms express more than the fact that the speaker approves or disapproves of whatever is being referred to: crucially, these adjectives are often associated with particular positive or negative *feelings* in the speaker. In order to remove the unwanted implication that the mental processes leading to speech are purely conscious and non-emotional, we can replace 'thought' in Ogden and Richards' diagram with the more neutral term 'psychology'.

Question Apart from emotion, what other aspects of psychology are relevant to the production and understanding of language? Which are of the most relevance to linguistic meaning?

The leftmost point of the triangle, the 'symbol', is the most straightforward. The symbol, in this terminology, is whatever perceptible token is chosen to express the speaker's intended meaning. In the case of spoken language, the symbols will be strings of speech sounds, in the case of written language, they will be marks on the page, and in the case of sign languages, they will be particular handsigns. Since in this book we are exclusively concerned with linguistic communication, we can replace the broader term 'symbol' with the simple 'language'.

The last apex of the triangle is the 'referent', or whatever things, events or situations in the world the language is *about*. Thus, the sentence *the dogs bark, the caravan goes by* has as its referent a particular situation: a situation in which certain dogs bark and a certain caravan goes by. Within that sentence, the expressions *the dogs* and *the caravan* also have referents: the actual dogs and caravan being spoken about. Note that someone who hears this sentence does not necessarily know what the exact referents of these nouns are; in the absence of any special knowledge about which dogs and caravans are being referred to, a hearer could only identify the dogs and caravan in question if the sentence was spoken when they were actually present (and even then they would have to assume that the hearer was talking about the dogs and caravan at hand, not some others).

This leads to the important point that we do not have any access to the world as it actually, objectively is. The only referents we can know are ones which are perceived by our senses or imagined in our minds: ones for

which, in other words, we have **mental representations** (see 1.6.2 below). The dogs and caravan in question are only available and known to us insofar as they can be **represented**, that is perceived, remembered, or otherwise thought about by us. The world of referents, that is, must be considered not as a world of real external entities, but as a world of representations which are *projected* by the mind. Another way of putting this would be to say that the world of referents is *within* the domain of psychology. As humans with minds, we have no access to *the* world, with a definite cast of fixed, pre-established referents. All we can know, and all that can be relevant to our understanding of language, is the world as it is represented by our minds through perception, memory, imagination or other experience. And since we are all different, the ways in which we perceive, remember or imagine referents are also likely to differ in some ways.

Question What problems might the existence of differing representations of the same referent pose for understanding meaning?

We can now consider the relations between the three points of the triangle. First, note that psychology has a causal relation to both referent and symbol. On the side of the symbol, the causal relation to psychology is explained by the fact that, as already observed, it is our minds that create language by choosing and constructing the particular linguistic expressions used. It is in our psychology that the decision to speak is made, and the particular words used are chosen. In the case of the referent (which, as we have already seen, must itself already be considered as within the domain of psychology), the causal relation comes from the fact that in using language we intend our words to have a certain referent. For example, if I point to a car parked on the street and say 'that car has its lights on' I intend my words to refer only to the car in question, and not to any of the others that also happen to be present. I have, in other words, chosen this car, rather than another, as the referent of my words, and I expect the hearer of my words to do the same.

In contrast to the causal relations on the psychology-symbol and psychology-referent sides of the triangle, there is no causal relation between symbol and referent. Words have no direct relation to the things they stand for. There is no inherent relation between a string of sounds and a particular referent: this is the reason that different languages use entirely different words for the same thing. The only reason *dogs* refers to dogs and *caravan* refers to a caravan is that these are the referents which English speakers have learnt to associate with them, and this is a fact about people's psychology rather than an essential connection between the words and the objects to which they refer. Even **onomatopoeic** words like the names for animals' calls (e.g. 'cuckoo', 'moo', 'quack' and 'meow'), which might be thought to constitute an exception to this rule, since their sounds are similar to the calls they represent, are not in fact any different. Even though there is certainly a similarity between word and referent, this similarity is a conventional one which, just as for other words, has to be learned (that is why different languages represent these sounds differently: for example, 'quack' in French is *coin-coin*). The connection

between onomatopoeic words and their referents is thus mediated by the psychology of language users.

In light of these remarks, we can redraw the semiotic triangle as in Figure 1.3:

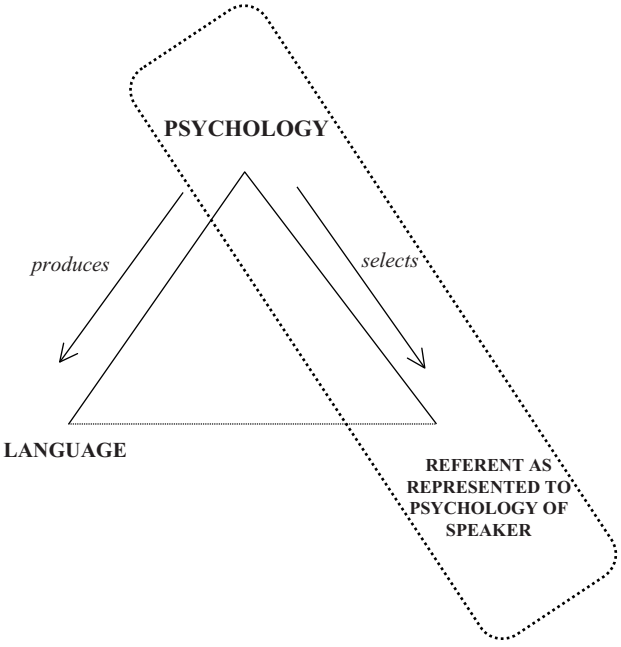

FIGURE 1.3
The semiotic triangle, re-labelled.

1.4 Some initial concepts

In this section we introduce some important concepts which we will need in the chapters that follow. The exposition here is only preliminary; each concept will receive a more detailed treatment later in the book.

1.4.1 Lexemes

To linguists and non-linguists alike, the *word* is the most basic and obvious unit of language. But in many languages, units which we would want to recognize as a single word can appear in many different morphological forms. Thus, in English, *go, goes, went, have gone* and *to go* are all forms of the verb *to go*. Other languages have many more morphological variants of a single word-form. In Ancient Greek, for example, a single verb, *tithēmi*, which means 'put', has several hundred different forms, which convey differences of person, number, tense and mood, such as *e-thē-ka* 'I put', *tithei-ētēn* 'you two might put', *thō-men* 'let us put', etc. But these different forms only alter some aspects of the meaning of the word. Both *go* and *tithēmi* share a large component of meaning between their different forms: *tithēmi* always has the sense 'put', and the forms of the verb *to go* always have the sense 'go', regardless of whether the sentence in question is 'I went' or 'you have gone'. For this reason, a semantic description does

not need to treat all the variant morphological forms of a single word separately. The **lexeme** is the name of the abstract unit which unites all the morphological variants of a single word. Thus, we can say that *go, goes, went, have gone* and *to go* all are instantiations of the lexeme *to go*, and *e-thē-ka*, *tithei-ētēn* and *thō-men* are all instantiations of the lexeme *tithēmi*. We usually refer to the lexeme as a whole using one of the morphological variants, the **citation form**. This differs from language to language: for verbs, for example, English, French and German all use the infinitive as the citation form (*to go, aller, gehen*), whereas Warlpiri uses the non-past form of the verb (*paka-rni* 'hitting', *yi-nyi* 'giving').

Not all languages have a word for 'word'

Not all languages have a word corresponding to English 'word': Warlpiri, again, makes no distinction between 'word', 'utterance', 'language' and 'story', all of which are translated by the noun *yimi*. In Cup'ik (Yup'ik, Central Alaska), the word for 'word' also means 'sayings, message' and 'Bible' (Woodbury 2002: 81). Dhegihan (Siouan, North America) has a single word, *íe*, referring to words, sentences and messages (Rankin *et al.* 2002).

1.4.2 Sense/reference/denotation/connotation

As we have already seen, the English word 'meaning' is rather vague. One important distinction we can make within the general notion of a lexeme's meaning is between its **sense** and its **referent** (or **reference**). To simplify the introduction of these terms, we will confine our discussion to nouns; we will see in 1.6.1 how they apply to other lexical categories.

The sense of a lexeme may be defined as the *general meaning* or the *concept* underlying the word. As a first approximation, we can describe this as what we usually think of as contained in a dictionary entry for the word in question, although we will see later that this characterization needs significant modification. The notion of sense can be made more explicit through contrast with the category of referent. A word's **referent** is the object which it stands for on a specific occasion of use. For example, consider (28):

(28) *The queen has fallen off the table.*

If I am talking about a rowdy evening at Buckingham Palace in 2009, the referent of the word *queen* is Her Majesty, Elizabeth II, and the referent of the word *table* is a particular piece of English royal furniture. But if I am talking not about Elizabeth II but about Queen Margrethe of Denmark, the words *queen* and *table* have different referents: not Elizabeth II and the English piece of furniture, but Margrethe and the Danish one. On each of the occasions (28) is uttered, there is one and only one referent of each word.

A word's referent, then, is the particular thing, person, place, etc. which an expression stands for on a particular occasion of use, and it changes each time the word is applied to a different object or situation in the world. (As we will see in Chapter 3, not all uses of nouns are referring, but we will ignore this for the moment.) By contrast, a word's sense does not change every time the word takes on a new referent. Regardless of whether the referent of *queen* is Elizabeth II or Margrethe, its sense is something like 'female reigning monarch'. This is not to say, however, that 'female reigning monarch' is the *only* sense of the word *queen*. Another sense of *queen is* 'second highest ranking piece in a game of chess'. This would be the sense involved if I uttered (28) while talking about a game of chess in the café, where *queen* would refer to a particular chess piece. Yet another sense of the word *queen* is 'third highest card in a suit, behind ace and king': this would be the sense involved if I uttered (28) in reference to a game of bridge at the kitchen table. In these two cases, *queen* does not only have two new different referents, the particular chess piece and the particular card, but two new different senses as well: 'second highest ranking piece in a game of chess' and 'third highest card in a suit, behind ace and king'. In *all* the utterances of (28), by contrast, 'table' has the single sense 'piece of furniture with raised flat surface used for putting things on, eating at, etc.'. Obviously, words like *queen* and *table* stand for many different people and objects in the world: they have, in other words, many different referents. The referents change each time we talk about a different queen, or a different table. The entire class of objects, etc., to which an expression correctly refers is called the expression's **denotation**.

Words have the referents they have by virtue of a certain act on the part of the speaker, which we will call the act of **reference**. We will use this term to describe what the speaker does in applying a particular language expression to a particular referent in the world. In uttering (29), for example,

(29) *Dr Schreber suffered his first illness in the autumn of 1884.*

the speaker *makes reference* to a certain person, Dr Schreber, to a certain disease, his first illness, and to a certain time, the autumn of 1884. These individual objects are the referents of the words in (29), and it is only in virtue of an act of reference, undertaken by the speaker, that the words 'Dr Schreber', 'first illness', and 'the autumn of 1884', have the referents they do. Since reference is an act, it is subject to exactly the same problems as all other human ventures, and it may not be successful. Thus, if I suddenly say to you 'I saw that cat again', and you don't know what cat I mean, reference will not have been successful. Even though I, as speaker, have referred to a particular cat, you (the hearer) are not able to **recover the referent** intended, i.e. identify the cat in question.

Reference, referents and denotation

Some writers use the term **reference** and **denotation** interchangeably, but in this book we will distinguish the two. An expression's denotation is the class of possible objects, situations, etc. to which the word can refer. The term *reference*, by contrast, has two uses:

- as the name of the act by which a speaker refers to a referent;
- as a synonym of *referent*, i.e. as the term for the object(s) to which an expression refers on a particular instance of use.

In this book, we will not try to distinguish these two senses of *reference* with separate terminology. *Reference* sometimes means the act of referring, and sometimes means a referent. The context will remove any doubt about which sense is intended.

Sense, reference and denotation are three aspects of what is commonly conveyed by the loose term 'meaning'. A fourth, very important aspect of meaning is **connotation**. Connotation names those aspects of meaning which do not affect a word's sense, reference or denotation, but which have to do with secondary factors such as its emotional force, its level of formality, its character as a euphemism, etc. 'Police officer' and 'cop', for example, have very different connotations, but similar denotations, as do the following pairs:

(30) *brat* and *child*
toilet and *rest room*
country town and *regional centre*
underprivileged area and *slum*
mutt and *dog*
doctor and *quack*
incident and *accident*

We will consider connotation again in Chapter 11.

QUESTION Think of some other pairs of words in English or any other language you know which have different connotations. Would you also want to say that they have different senses?

1.4.3 Compositionality

All human languages have the property of **productivity**. This is simply the fact that the vocabulary of any given language can be used to construct a theoretically infinite number of sentences (not all of which will be meaningful), by varying the ways in which the words are combined. For example, given the words *the, a, has, eaten, seen, passing, contemporary, novelist* and

buffalo, the following figure among the large number of meaningful sentences that can be constructed:

(31) *The novelist has seen the buffalo.*
A novelist has eaten the buffalo.
A contemporary novelist has seen a buffalo.
The novelist has seen a passing buffalo.
A buffalo has eaten a passing contemporary novelist

and so on. (We can also construct ungrammatical sentences like *A the novelist eaten passing has*, but since these are meaningless we will ignore them here.) Most people have probably never heard (32) before:

(32) *There are no remains of ancient Indian aircraft technology*

yet, as speakers of English, we understand immediately what it means. How does this ability arise? One answer is that meaning is **compositional**. This is to say that the meanings of sentences are made up, or **composed**, of the meanings of their constituent lexemes. We understand novel sentences because we understand the meanings of the words out of which they are constructed. Since we know the individual meanings of *there, are, no, remains*, *of*, *Indian,* and so on, we know the meaning of any grammatical sentence in which they are combined. On the contrary, if a novel sentence contains a word which we do not know, we do not know what the sentence means. Thus, if you are told that *the distribution of seats was aleatory*, and you do not know that *aleatory* means 'random', then the sentence, taken as a whole, will not be meaningful. It is important to note that not all combinations of words are necessarily compositional. One especially important category of non-compositional phrase is **idioms**. For example, if I say that so-and-so has *thrown in the towel*, most English speakers will recognize that I am not talking about anyone literally 'throwing' a 'towel', but that I simply mean that the person in question has given up on whatever venture is being spoken about. The phrase *throw in the towel*, then, is *not* compositional, since its overall meaning, 'to give up', does not derive from the meanings of its individual component lexemes.

QUESTION In the following sentences, which of the highlighted expressions can be considered compositional, and which are idioms? Do any belong to some third category?

If you ***keep on*** *making that noise I'll* ***go through the roof.***
He's just ***kicked the bucket.***
Stop ***dragging the chain****: we'll never get there.*
We've run out of time, so we'll have to ***wrap things up****.*
Can you ***run off*** *twenty more copies?*
After the delay the plane ***took off*** *as normal.*
I'll ***take*** *twenty per cent* ***off*** *the price.*
This is a ***nice and hot*** *cup of tea.*
My hands are ***lovely and warm****.*
Try and get a better deal next time.
Hello down there!

Based on the distinction between the meanings of words and the meanings of sentences, we can recognize two main divisions in the study of semantics: **lexical semantics** and **phrasal semantics**. Lexical semantics is the study of **word meaning**, whereas phrasal semantics is the study of the principles which govern the construction of the meaning of phrases and of **sentence meaning** out of compositional combinations of individual lexemes.

1.4.4 Levels of meaning

The distinction between word meaning and sentence meaning, then, defines a basic contrast between lexical and phrasal semantics. Another important contrast is the one between **sentence meaning** as just described and **utterance meaning**. We can define sentence meaning as the compositional meaning of the sentence as constructed out of the meanings of its individual component lexemes. But the meaning of a sentence as built up out of its component parts is often quite different from the meaning it actually has in a particular context. In everyday talk we regularly use words and expressions ironically, metaphorically, insincerely, and in other 'non-literal' ways. Whether there is any principled theoretical difference between these non-literal ways of talking and the literal ones, and, if so, what it is, is an important question which we will discuss in Chapter 7; for the moment, we can simply recognize that there are many uses in which words seem to acquire a strongly different meaning from the one they normally have. Suppose that while cooking Peter has just spilled a large quantity of spaghetti carbonara all over the kitchen floor. Hearing the commotion, Brenda comes into the kitchen, sees what has happened, and utters (33)

(33) *You're a very tidy cook, I see.*

It is clear that Brenda doesn't literally mean that Peter is a tidy cook, but that she is speaking ironically. What she actually means is the opposite of (33): Brenda is drawing attention to the fact that Peter has precisely *not* been a tidy cook. In cases like this, we say that there is a difference between sentence meaning and utterance meaning. The sentence meaning of (33) is the literal, compositional meaning as built up from the meanings of the individual words of the sentence. If we did not speak English, we could discover the sentence meaning of (33) by finding out what its translation was in our own language. The **utterance meaning**, by contrast, is the meaning which the words have on a particular occasion of use in the particular context in which they occur. (Utterance meaning is sometimes referred to in other books as **speaker meaning**. But since the role of the hearer is just as important as that of the speaker, the more neutral term **utterance meaning** is preferred here.) The utterance meaning is the one which is picked up in the conversation. In reply to (33), Peter might well say (34):

(34) *I'm sorry. I don't know how I could have been so clumsy.*

But if Brenda's comment in (33) was meant literally, the reply in (34) would be very strange: people do not usually have to apologise for being

tidy. What (34) shows is that it is the utterance meaning, not the sentence meaning of (33) to which Peter is reacting: given the situation, Brenda is clearly not congratulating him on his tidiness as a cook, and it is the utterance meaning which forms the basis for the continuation of the conversation.

The distinction between sentence meaning and utterance meaning is also linked to the difference between **semantics** and **pragmatics**. For those linguists who accept such a division, semantics is taken to study sentence meaning, whereas pragmatics studies utterance meaning and other principles of language use. The job of semantics is to study the basic, literal meanings of words as considered principally as parts of a language system, whereas pragmatics concentrates on the ways in which these basic meanings are used in practice, including such topics as the ways in which different expressions are assigned referents in different contexts, and the differing (ironic, metaphorical, etc.) uses to which language is put. As we have already seen, a division between semantics and pragmatics is by no means universally accepted in linguistics. Many 'pragmatic' topics are of central importance to the study of meaning, and in this book we will not recognize any absolute distinction between the two domains.

1.5 Object language and metalanguage

Like any other branch of linguistics, semantics deals with the words, phrases and sentences with which we communicate. But for semantics the immediate objects of study are not these words, phrases and sentences themselves, in the sense of the sounds, sequences of letters or handsigns which we utter or perform and can then write down or record. As the study of meaning, semantics is interested in something which cannot be perceived directly through our senses, but which, in one way or another, we *experience* in using and thinking about language. We cannot see, hear or touch a word's meaning: meanings are things we understand. It is not meanings that go between speaker and hearer: the only things that are transferred from one speaker to the other are sound waves in the air. This means that in order to get started in semantics, we need a way of identifying meanings and bringing them to light in an unambiguous way so that we can begin to study them.

The main way in which we normally reveal the meanings of linguistic expressions is, quite simply, by describing them in language. But since it is language that we're interested in in the first place, we need to distinguish between the language *whose* meanings we want to describe and the language *in which* we couch the descriptions. The language whose meanings we are describing is called the **object language**. The language in which we describe these meanings is called the **metalanguage**.

When we propose a metalanguage description of the meaning of an object language expression, we are using one type of meaning (the

meaning of the metalanguage expression) to explain another kind (the meaning of the object language expression). Let us take the example of English as the metalanguage for Dutch, as would be the case if we were explaining the meaning of the Dutch word *groot* to an English speaker. One possible metalanguage explanation, or **definition**, that we could give of *groot* in English would be the word 'tall', as in (35):

(35) *Dirk is groot, maar Lou is klein.*
'Dirk is tall, but Lou is short.'

The object language expression *groot* is here defined by the metalanguage expression 'tall'. But note that not all aspects of the word 'tall' are relevant to this definition: it is completely irrelevant to the definition of *groot* that the metalanguage definition we have chosen, 'tall', has four letters, or is a monosyllable, or starts with the consonant /t/. All these phonetic and orthographic details are irrelevant to semantics, since the only thing that matters for the purpose of defining *groot* is what 'tall' *means*. To tell someone that *groot* means 'tall' is to make a statement about two meanings, and to say that these two meanings are the same. The particular phonetic and other characteristics of the metalanguage term are therefore irrelevant: as long as the person for whom the definition is intended understands the meaning of 'tall' in English, the definition is successful.

This confronts us with an interesting problem. For couldn't it be objected that, in one way, we haven't actually explained anything when we define *groot* as 'tall'? We have certainly given a definition of the word which will help an English speaker to understand the meaning of the Dutch sentence. But if we want to go beyond the problem of allowing people to translate from one language to another, hasn't our analysis left something crucial out? On hearing our explanation that the meaning of *groot* is 'tall', someone might easily object by pointing out that this explanation only shows an equivalence between two words in English and Dutch, and does nothing to explain what this meaning, which both *groot* and *tall* express, actually *is*. 'I know what *groot* means in English', they might say, 'but you haven't told me what it actually is for something to mean something.' And even though we could go on to give a more detailed explanation of 'tall', perhaps using terms like 'elevated in stature', 'not short', etc., for as long as we continue to explain the meaning of a word by using the meanings of other words, we will not have satisfied our objecter's curiosity.

For many linguists, this objection is rather forceful. As long as we go on defining meanings by other meanings, we leave out the essential task of explaining what meaning actually is. We can see this very clearly by considering the case of **circular definitions**. Consider someone who wants to find out the meaning of the English word 'humorous'. One possible definition of 'humorous' would be 'droll'. But this definition would only be effective if the meaning of 'droll' was already known. If it was not, it too would need to be explained: 'droll', perhaps, could be plausibly explained

through the metalanguage definition 'amusing'. 'Amusing', in turn, could be defined as 'funny', as in (36).

(36) humorous ⇒ droll ⇒ amusing ⇒ funny.

Depending on the person for whom the definition was intended, this chain of definitions would sooner or later achieve its purpose: if the person knew the meaning of 'funny', we could stop the explanation at this point, so that 'humorous' would have been defined through 'droll', 'droll' through 'amusing', and 'amusing' through 'funny'. It is obvious, however, that this chain could not go on for ever. Sooner or later we would run out of new words: if the language learner did not know even what 'funny' meant, we can imagine giving up in frustration, and saying, simply, '*funny* just means "humorous"'. In this case, it's clear that our unfortunate language learner would be none the wiser, since 'humorous' was the word whose meaning was originally in question. Since 'humorous' has been used both as an object language term and a metalanguage term, the definition is **circular** and does not succeed in telling us anything new:

(37)
humorous	⇒	droll
⇑		⇓
funny	⇐	amusing

Clearly, then, for as long as we remain within the circle of definitions by substituting one word or phrase as the definition of another, we remain confined within language. The lexical resources of any language are limited: at some point, the metalanguage definitions will *have* to include object language terms, and thereby introduce circularity. We can continue to refine our definitions and search out the most precise and explanatory ways of couching them, but in contenting ourselves with this task we will not have provided any account of what the meanings we are defining actually are, nor of how they relate to any of the three points of the semiotic triangle. In particular, we will have left it completely obscure what it is for a speaker to understand the meaning of a word. If I understand the meaning of 'droll', then the definitional chain can be stopped. But what does it mean to say that I understand the meaning of 'droll'? What is it that I actually understand? For many linguists, the fact that we cannot answer these questions about meaning by remaining inside the definitional circle means that we have to look outside language for answers. If linguistics is to play a part in explaining the way language can be actually used by real speakers, we need to find a point at which the circle can be broken in order to link meaning in with something non-linguistic. We will consider a few proposals about how this could be done in the next section.

1.6 Breaking the circle

As pointed out by Quine (1961: 47), until the development of 'a satisfactory explanation of the notion of meaning, linguists in semantic fields are in

the situation of not knowing what they are talking about'. This is perhaps not such a dire situation as it sounds: after all, empirical investigation always aims to increase our knowledge of some unknown phenomenon, provisionally characterized using ordinary language. As the inquiry proceeds, we get a sharper idea of the nature of the thing being studied, and it may not matter that in early stages we have to rely on notions for which we cannot yet give any satisfactory explanation. Many fields of empirical inquiry begin with only hazy and imprecise conceptions of the real object of their investigation. The history of genetics is a case in point. Mendel, acknowledged by most historians as the founder of the field, discovered the principles of inheritance without any understanding of either chromosomes or DNA, both of which later became central parts of the theory of cell biology. The fact that his advances were thus made in ignorance of the fundamental nature of inheritance does not in any sense discredit them: Mendel might not have known exactly what his discoveries were ultimately about, or what the mechanisms were that implemented the facts he observed, but his rigorous investigations meant that he was able to reach valuable conclusions which would only be fully characterized later. The fact that he could not have precisely characterized the nature of the phenomenon he was observing was not an obstacle to progress (see Gribbin 2002: 536–541 for discussion).

Still, to say the least, it would obviously be *useful* if we had some initial idea about what meaning is best thought of as being – of how, in other words, we can break the definitional circle. This preliminary idea will help us to formulate the best set of specific questions to ask in our investigation. In this section, we will consider several suggestions about how the definitional circle might be broken and the notion of meaning explicated in a way which might satisfy objections like Quine's.

1.6.1 Meanings as referents/denotations

One way to break the definitional circle would be to stress the role of the referent or denotation as the main component of the meaning of a linguistic expression. Under this theory, metalanguage explanations of a meaning should be seen as names of the referents of the object language term. As we saw in Section 1.2.1, ordinary discourse about language in English often seems to make an implicit identification between an expression's meaning and its referent:

(38) *In Sydney, 'the bridge' means the Harbour Bridge.*

The 'meaning' of 'bridge', the speaker of (38) seems to be suggesting, is the actual harbour bridge itself. 'Bridge', we might say, *means* what it *refers to*; its meaning on any one occasion of use *is* its referent. Outside of the narrow context of (38), we could say that the meaning of *bridge* in general is just its denotation – the class of all bridges. This identification of meaning and referent/denotation succeeds in breaking the circle because it identifies meaning with non-linguistic objects in the world: the meaning of 'bridge' on a particular instance of use is the real bolts and metal structure. Given

this interpretation of the meaning of 'bridge', it doesn't matter that we would eventually run out of new words with which to define it, since we can ultimately analyse its meaning **ostensively**, i.e. simply by pointing at the actual bridge itself.

As was pointed out in the discussion of the semiotic triangle (1.2.1), the referents of expressions must be taken not as actual objects in the world, but as representations in the world as projected by the speaker. This means that in order to understand reference we already have to invoke the realm of speakers' individual psychologies, the particular 'versions' of the world as projected by their psychology. The postulation of the world of projected representations allows us to avoid an objection which might otherwise count against referential theories of meaning. This is the objection that it is often the case that there simply is no referent for a given expression, as in (39a–c), or that the referent is unknown, as in (39d–f):

(39) a. *The German victory in World War II*
b. *Robin Hood's private helicopter*
c. *The water on the Moon*
d. *The most distant point from Earth*
e. *The first person ever to use language*
f. *The fate of the environment*

A theory which identified meanings with *real world* referents would have to say that the expressions in (39a–c) simply have no meaning, since the things they refer to never actually existed, or are impossible; and it would have to say that the meaning (referent) of the expressions in (39d–f) was unknown, since although we can be confident that all of the things referred to by the expressions exist, we do not know what they are. But if referents are taken to be representations projected within the realm of people's psychology rather than real objects in the actual world, this problem disappears. Whether or not there is any object referred to by the words *Robin Hood's private helicopter*, we can easily think of situations in which a speaker might simply imagine, pretend or otherwise entertain the possibility that such a helicopter did exist. For the speaker of (39b), then, the referent of *Robin Hood's private helicopter* can be taken as the speaker's representation of the helicopter in their projected world. The reader will easily see that similar explanations can be constructed for the other examples in (39).

The identification between meaning and reference may be successful in breaking the definitional circle, but it leads to a very fragmented picture of the nature of language: on the reference theory of meaning, 'bridge' has as many different meanings as it has different referents. This variety clashes with our pretheoretical intuition that the meaning of *bridge* is actually something much more unitary: although there are many different individual bridges out there in the world, the meaning of the word *bridge,* or, we might say, the *concept* of a bridge is a unified, single entity.

The idea that an expression's meaning is its referent is at least easy to understand for nouns referring to discrete, concrete things. But it is much

less clear what the referents of other lexical categories might be. What are the referents of abstract nouns like *scandal*, *generosity* or *impermanence*? Since there is no isolable object in the world to which these nouns apply, the notion of a referent is rather hard to invoke. And what about adjectives like *sweet*, *polished* or *ineffectual*, or verbs like *to have*, *to allow* or *to go*? In the case of 'grammatical' words the problem is even greater: what is the denotation of *of*, or of *the*? These cases all pose problems for the referential theory of meaning: because the words have no referents/denotations, they are left without any specifiable meaning. Yet it is obviously the case that these words *do* have meanings, which we can paraphrase metalinguistically and explain to others. We will consider this question further in Chapter 6.

A second problem with the theory of meaning as reference is the fact that a single referent may often be referred to by a variety of different expressions. Thus, the expressions in the two halves of (40a–d) each pick out just a single individual:

(40) a. *The first country to adopt a law requiring parental leave; the home country of IKEA*
b. *The most frequently handed in, and the least frequently claimed, object on the Tokyo subway; portable device with handle used for protection against rain*
c. *The inventor of Chupa Chups; friend of Salvador Dali and husband of Nuria Serra*
d. *Institution for lending money; institution for depositing money*

In (40a) we have alternative ways of referring to Sweden, in (40b) of umbrellas, in (40c) of the Spanish confectionery king Enric Bernat Fontlladosa, and in (40d) of the word *bank*. Yet we surely do not want to say that the meanings of these expressions are the same. While the objects referred to by the expressions 'institution for lending money' and 'institution for depositing money' have the same denotation – banks – they clearly don't have the same sense. We could imagine a bank which suddenly stopped lending money even though it continued to accept deposits: something like this, indeed, happened during the Argentinian financial crisis of 2002 and the global one of 2008. If meaning simply *is* reference/denotation, then examples like this should not be possible. The fact that linguistic expressions can be identical in reference but different in meaning leaves us no choice but to conclude that there is more to meaning than reference/denotation.

1.6.2 Meanings as concepts/mental representations

The referential/denotational theory of meaning broke the definitional circle by emphasizing the referent side of the sense/referent pair. Another way out of the circle is to identify meanings with **concepts**: the metalanguage definitions of an object language meaning, in this theory, are the names of the concepts associated with the object language term. The use of the term 'concept' in linguistics derives from philosophy, where it has

a very long history of discussion and controversy. For our purposes, concepts can be seen as a way of talking about the basic constituents of thought. In the words of Prinz (2002: 1) '[w]ithout concepts, there would be no thoughts. Concepts are the basic timber of our mental lives.' As we will see later, many investigators think it is necessary to distinguish between **primitive** concepts and others. On this view, our stock of concepts is built up from a stock of primitive concepts, which cannot themselves be broken down into any constituent parts. This level of primitive concepts is the bedrock of the whole conceptual system; all other concepts can be analysed into combinations of these simpler primitives, just as all molecules can be analysed down into their basic component atoms. For the moment, we will not distinguish between primitive and non-primitive concepts; we discuss the distinction in detail in Chapters 2 and 8.

If we imagine the process of thinking as a sort of internal conversation with ourselves, then concepts are the individual words and expressions of which this conversation consists. Concepts are implicated in practically every aspect of our mental lives. It is on the basis of concepts that we determine things' identity: if I want to know whether some animal is a mammal or a marsupial, for example, I subconsciously compare its properties against the properties of the concepts MAMMAL and of MARSUPIAL. Concepts are also needed to explain how we recognize objects in the world as themselves: if I know, when looking at a golf ball, that it is a golf ball, it is because the visual image accords with my concept GOLF BALL. Similarly, it is because of the involvement of concepts that our thought has continuity: if I am studying semantics, for example, I am progressively refining concepts like MEANING and REFERENCE with which I understand the functioning of language, and it is the same concepts MEANING and REFERENCE which are developed over the entire time I am studying. We have concepts corresponding to abstract words like *democracy*, *possession* or *time*, but equally for everyday ones like *hand*, *red*, *go*, *hungry*, *anticlockwise* and *up*.

One very common way of describing language in the Western tradition, going back to Aristotle, is to see language as *communicating ideas*: on this understanding, we choose the particular words we use in order to achieve the closest fit with the particular ideas we have. And, indeed, as pointed out by Reddy (1993), we often talk, in English and many other European languages, as though language was a receptacle into which we put ideas in order to transfer them to the hearer, as in (41):

(41) *There are a lot of ideas* in *that sentence.*
You can get *the same meaning* across *in different ways.*
I can put *the same* idea in *different words.*

Language, then, is often spoken about as though it was the 'conduit' for ideas. A natural extension of this common understanding of language is that what words actually mean are ideas or concepts. Thus, the meaning of the word 'tolerant' is our concept TOLERANCE: when we say 'Oliver is

tolerant', we are attributing to Oliver certain properties which together define our concept TOLERANCE, like patience, kindness, respect for the opinions of others, and so on. These properties can be thought of as combined together into the concept TOLERANCE, rather like the different components of a definition of *tolerance* in a dictionary.

The hypothesis that meanings are concepts has considerable attraction. First, it answers to the intuition that language is intimately connected with the rest of our mental lives. It does seem precisely to be *because of* the thoughts and concepts we have that we use the words we use. If I say 'horse-drawn carriages are old-fashioned', then this will often be because this is exactly what I think: I am reporting a link between the concepts HORSE-DRAWN CARRIAGE and the concept OLD-FASHIONED. Language and thought are very hard to tease apart: whether or not we always think 'in language', we often need to use language to externalize the results of our thought, to bring these results into the public domain for the purposes of communication, and it seems to be in language that most of our ideas can be given their most precise form. Since there is this clear causal connection between language and thought, the idea that the meanings expressed through language correspond to concepts is a neat way of effecting the link between the world of public, external communication and our private, mental lives.

Second, the conceptual theory of meaning has often been taken to explain compositionality and relations between meaning. The concept HORSE-DRAWN CARRIAGE can be seen as built up from the concepts HORSE and the concept CARRIAGE, as well as some third element corresponding to the word 'drawn'. Similarly, the meaning of the linguistic expression *horse-drawn carriage* has these very three elements (at least), and they can be individually changed to create different expressions with different meanings. In such cases, we can explain the changed meanings as corresponding to changed concepts. Thus, instead of a horse-drawn carriage, we can imagine an *ox-drawn carriage* or a *horse-drawn plough*: in these cases, we have substituted the concepts OX and PLOUGH for HORSE and CARRIAGE, and these substitutions explain the altered meaning of the expressions. The conceptual hypothesis also explains certain other links between the words 'horse drawn carriage' and other words. For example, a little reflection will reveal that HORSE-DRAWN CARRIAGE is a member of the more inclusive concept MEANS OF TRANSPORT, and is linked, by association, with such concepts as COACHMAN, PASSENGER, REINS, WHEEL, etc. It is these conceptual links which ultimately explain the comprehensibility of sentences like (42a) and (43a), and of how they are different from those of (42b) and (43b):

(42) a. *A horse-drawn carriage is an old-fashioned means of transport.*
b. *A horse-drawn carriage is an old-fashioned cheese.*

(43) a. *The coachman jumped down from the horse-drawn carriage.*
b. *The sunrise jumped down from the horse-drawn carriage.*

The meaning of (42a) and (43a) is clear and easily understood because the words all express related concepts. But since the concepts expressed as the

meanings of the words in (42b) and (43b) are not inherently connected, the meaning of these sentences is much harder to interpret.

Meaning relations like **synonymy** (sameness of meaning) are also easily explained by the conceptual hypothesis. Two words are synonyms if they have the same meaning. And 'having the same meaning' means 'instantiating the same concept'. Thus, 'Islamic' and 'Muslim' might be said to be synonyms, because the corresponding concept, which we can either refer to as MUSLIM or ISLAMIC, is identical.

Third, the hypothesis that meanings are concepts guarantees the genuineness of communication. Because meanings of words are concepts, two people who talk, agree or disagree about something are doing more than 'playing with words'; they are talking, agreeing or disagreeing about certain concepts, which are being compared and progressively reconciled with each other during the exchange. And as the concepts are complicated, easy, familiar or unfamiliar, so are the meanings. It is therefore the level of concepts that guarantees that genuine communication between people can actually take place.

What form do concepts take psychologically? This is an extremely controversial question. An answer favoured by many linguists, adopted from philosophy and cognitive science, is that concepts have the form of symbolic **mental representations**. Mental representations are the fixed mental symbols – the 'language of thought' – which are instantiated in our minds in some stable, finite medium, and which our thought consists in. On the view of concepts as mental representations, thinking and expressing meaning are both to be understood as the manipulation of mental symbols, in much the same way that using language is the manipulation of a fixed series of linguistic symbols in the medium of air, paper or handsigns. Communication, then, involves using the conventional names for individual mental representations. Since these individual mental representations belong to a language-like format in which the contents of mental events are expressed or recorded in the mind, their 'translation' into the words of natural language follows readily.

There are, however, a number of reasons we should be cautious in the claim that meanings correspond to concepts. We will mention only three now. First, some words seem more naturally compatible than others with an interpretation of their meanings as concepts. Thus, while it seems quite plausible to say that the meanings of *democracy, punctuation, panorama,* or *love* are concepts, this move is less obvious for words like *ouch*!, *me, you* or *this*, or so-called 'function' words, like *if, not, like* or *very.* Words like these do not seem to be able to call up the rich range of associations and inherent connections which characterize *democracy, love*, etc. The point here is not to rule out the possibility that the meaning of all these words may in fact correspond to concepts, but simply to suggest that the initial intuitive plausibility of this is not as great.

QUESTION Can you propose any 'conceptual' content for the above words? What about words like *brown*, *zig-zag* or *bitter*? If so, what is it? If not, why not?

Second, just like meanings, concepts cannot be seen or otherwise identified unambiguously. This means that their postulation is not immediately controllable by objective, empirical means. Psychologists and psycholinguists have certainly developed experiments in which the properties of particular hypothetical concepts can be experimentally tested, but this has only been done for a fraction of words, largely from well-known languages, and, like any experimental result, the conclusions are open to a variety of interpretations. It is therefore unclear, given the present state of research, whether the postulation of concepts is scientifically justifiable, or whether it is simply a term we have adopted from our untutored, pretheoretical views about the nature of our mental lives. There is, of course, no *in principle* problem with postulating unobserved entities in semantics – any science works by postulating the existence of unobserved (and sometimes unobserv*able*) factors which are hypothesized to explain the observed facts. It is simply that, in linguistics, the detailed experimental work is only starting to be done that would put these unobserved entities on a more solid empirical footing.

Third, even if an expression's meaning can partly be identified with the concept it evokes, there must be more to it than that. For example, if I say the words *Wallace Stevens was a poet and an insurance broker*, I do not mean that my *concept* of Wallace Stevens was a poet and an insurance broker: I mean that a certain real person, Wallace Stevens himself, was. Part of what I mean, then, is the actual, real-world person that my words refer to. And this real-world person could prove to have quite different properties from the ones reflected in my concept of him. For example, I might mistakenly believe that Wallace Stevens is the author of *Death of a Salesman*: that fact, then, forms part of my concept of Wallace Stevens. But this doesn't mean that when I say *Wallace Stevens was a poet and an insurance broker* I am saying something false, even though it isn't true that the author of *Death of a Salesman* was a poet and an insurance broker. What makes my words true does not depend on the concept I have of Wallace Stevens, but on who this expression refers to. This isn't just the case with proper names. Imagine that I'm confused about the difference between lyrebirds and bowerbirds. I can tell the two apart, but I wrongly believe that lyrebirds are the birds that decorate their nests, and that bowerbirds are the birds that are incredibly good mimics. When I tell someone that *the bowerbird has just come back*, my meaning isn't just that 'the bird that is an incredibly good mimic' has just come back; it's that *that particular bird* (whatever it's actually called) has just come back. So while we might want to say that words express certain concepts, there does seem to be an important referential component to meaning which goes beyond concepts.

The hypothesis that meanings correspond to concepts has been very popular in linguistics. For many semanticists, this hypothesis is not, as it might appear, an alternative to an identification of meaning with reference or denotation, but is rather complementary to it. This is because under the conceptual theory of meaning the semanticist's task is not simply over once the referents and denotations have been identified for

the words under investigation. Concepts can be identified with senses, the general meanings of words as considered separately from their specific reference on any given occasion of use. Thus, once we have identified the referents and so the denotation of the noun *fire*, we can go on to explore the features of our *concept* FIRE which may be relevant to language. These features go beyond the mere identification of the objects denoted by the word. For example, we will discover that there is a close link between our concept FIRE and such other concepts as HOT, FLICKERING, DANGEROUS, BURN, RED etc. These conceptual links are useful for three reasons. First, they explain the compatibility between the word *fire* and words like *hot*, *flickering*, *dangerous* and *burn*, in just the same way as for examples (42a) and (43a) above, and account for the fact that these words will often occur in similar contexts in actual language. Second, the conceptual theory can explain certain *extended meanings*, such as that some hot things like intense summer weather or spicy food may also be described with the adjective 'fiery': presumably this has something to do with the close conceptual link between our concepts HOT and FIRE. Last, and most important, the postulation of the concept FIRE as the meaning of *fire* explains why *fire* has the referents it has. Thus, to the question 'why are these things, and not different ones, called *fires*?', the conceptual theory of meaning gives the reply 'because only these objects, and not others, accord with the concept FIRE which the word *fire* expresses'. Clearly, these are extremely informal explanations. Nevertheless, the only reason that even this low level of explanatory depth is possible is the presumed link between language and concepts. If we could analyse the meaning of *fire* no further than by itemizing a list of its referents, none of these commonsense observations about the relation of *fire* to other words would be justified. The conceptual theory of meaning thus provides a convenient rationale for a fruitful investigative practice, and justifies many commonsense observations about meaning.

QUESTION How might concepts provide an answer to some of the problems of the referential/denotational theory of meaning?

1.6.3 Meanings as brain states

A natural thought about meaning is to identify it with brain states: understanding or intending a certain meaning, on this identification, would just be having the neurons of one's brain in a particular configuration. From one point of view, this identification seems very plausible. After all, isn't language ultimately a product of our brain? If we understood how the brain works, wouldn't we understand all the details of language, semantic ones included? Isn't it only the rudimentary state of our current understanding of brain processes that prevents us from giving the details of this identification? According to this line of thinking, semantics, along with the rest of linguistics, will one day be **reduced** to, or unified with, brain science, in the same way that the classical theory of genetics can be reduced to or unified with that of molecular biology. In other words, once brain science has progressed, we will no longer need the technical vocabulary of semantics, but

will be able to talk wholly in terms of synaptic connections, neurotransmitters, proteins and so on. In just the same way, the modern language of molecular biology, involving chromosomes, nucleotides and DNA sequences, has at least partly replaced the older one of genes as the best way of describing the details of inheritance.

This is an attractive position in many ways: brain states must ultimately cause all behaviour, including language. But it is going too fast to conclude from this that we will eventually be able to *identify* meanings with brain states and reduce semantics to brain science. The first obstacle is that it's hard to see how brain states could have the properties that meanings have. Meanings, the way we normally think of them, have mutual connections to each other – of synonymy, antonymy, class inclusion and so on. For example, the meaning 'cat' has the following relations with other meanings:

- it is an instance of a broader class of meanings, 'mammals', 'domestic animals', 'four-legged animals' and so on;
- it is, in some sense, the 'opposite' of the meaning 'dog';
- it can be synonymous with the meaning 'feline'.

These are facts about the meaning of *cat* that we will presumably want a theory of semantics to reflect. It is not clear, though, how this could happen in a theory which identified meanings and brain states. How can one brain state be the opposite of, or synonymous with, another? The brain is just physical matter; it makes as little sense to say that a state of brain matter is the opposite of, or synonymous with, another as it does to say the same of the state of the electrons in my computer at any one time. It therefore seems that meanings have a property which prevents them from being completely identified with brain states.

This problem is just one instance of the broader problem of **intentionality**. Intentionality is the term philosophers use to describe the essential *aboutness* or *contentful nature* of language. A word like *cat* has a certain psychological content: it refers to (is about) a certain class of creatures, cats. The same is true of a verb like *sit*: this refers to, or is about, a certain class of actions, the actions of sitting. Many philosophers think there is something deeply special about intentionality, in that it is a property that is distinctively mental: purely physical things like my brain or my computer, which consist of configurations of electrons, just aren't the types of thing which can possess intentionality. Electrons, whether in my brain or in my computer, aren't about anything; they're just *there*. As a result, any attempt to simply identify something intentional like language with something non-intentional like a brain state cannot be successful.

How do we square this with the obvious truth that it is the brain that is ultimately responsible for linguistic production and understanding? If meaning is one of the factors to be taken into account in the production of utterances, and if brain processes will ultimately explain the *whole* production of utterances, then surely they must explain meaning too! It would just be illogical to say that everything that happens in language is determined by brain processes, and in the same breath to exclude

meaning. Here we need to invoke the concept of **levels of explanation** or **levels of description**, an important notion in cognitive science discussed by Marr (1982). Attending to the notion of levels of explanation/description will show us that there is room for *both* intentional meanings *and* non-intentional brain states in our explanations of language.

Consider a computer chess program. There seem to be several levels on which we can describe and explain what the program is doing. The first is the informal level, on which we describe the computer as simply following the rules of chess with the aim of beating its opponent. A particular move might be described as 'castling to protect the king', for example, or 'taking a pawn', or 'sacrificing the bishop'. This mode of description uses ordinary, everyday vocabulary of the sort we could also use for explaining people's behaviour. It makes reference to beliefs, intentions, and desires: the computer's belief that its king could be threatened, its desire to protect it, and its intention to castle in order to achieve this. In one way, the computer doesn't really have beliefs, desires or intentions, of course, but we talk as though it does, since this is a useful way of describing and understanding what is happening. Let's call this the **intentional** level of explanation.

A second, lower level of explanation is more detailed: this level consists in specifying the different steps of the program that the computer is running. This explanation wouldn't use the ordinary terms of everyday intentional, psychological explanation, but would lay out the sequence of individual input and outputs that the computer processes. The computer has some way of representing the input positions (the position of the pieces after each move), and a set of algorithms for turning inputs into outputs in a way that makes it likely to win. Let's call this the **algorithmic** level of explanation. Understanding this level will give us a detailed way of predicting what the computer is going to do next: if we have access to the specific program it is running, we can work out its next move in advance. Notice that there are several different ways in which the intentionally described actions the computer performs could be realized algorithmically. There's more than one possible chess program that a computer could run, yet all of them produce behaviour which is open to a single type of intentional explanation: the difference between different chess programs disappears at the intentional level of explanation where, whatever program the computer is actually running, we can still always describe it as 'castling to protect the king', 'taking a pawn', 'sacrificing the bishop' and so on. The details of the program become invisible as we move to the higher level.

Finally, there's the lowest level, the level of **implementation**: this is the level of description/explanation which concerns the specific way in which the algorithm is instantiated physically in the particular machine involved. Just as a single fact on the topmost intentional level can correspond to several different states on the lower algorithmic level, so a single algorithm can be implemented in multiple ways in an actual physical machine. This is most obvious if we think about the difference between the most up-to-date type of computer, which runs on a solid-state drive, a

conventional one using a spinning hard disk, and an old-fashioned one using magnetic tape or punch cards. All these machines can run the same algorithms, but the physical details of how they do so are completely different.

Clearly, all three levels of explanation are necessary to understand what is going on when the computer plays chess. Since there is a variety of possible physical realizations of the program on the implementational level, the next highest level, the algorithmic one, gives us a powerful way of abstracting from the details of the actual physical system that is performing the operations and describing the inputs and outputs of the program. But it is the intentional level that is the most relevant when we ask *why* the computer is behaving as it is. The intentional level, which consists of explanations like 'protecting the king', 'taking a pawn', 'sacrificing the bishop', makes sense of the computer's actions as a chess-player, not just as a machine. The algorithms and their physical instantiations are just a meaningless sequence of actions if we can't place them in the context that allows them to make sense, and it is only the intentional level that does this. Marr (1982) draws an analogy with trying to understand bird flight: we can't understand bird flight by limiting ourselves to feathers, the implementational level. We have to go beyond this to look at the wider place of feathers within a complex of notions like lift, air pressure, energy, gravity, weight and so on. Studying the physical constitution of the feathers in the absence of these other considerations will be fruitless.

Language is arguably the same way. Studying brain states will only tell us how language is implemented. It will tell us nothing about the higher-level relations that tie this implementation in with the rest of our psychology. As a result, meanings are unavoidable as part of the explanation of utterances. If I tell you that *my head is killing me*, then part of the explanation for my utterance involves my belief that my head hurts, my desire to communicate this fact to you, and the fact that those words convey that idea *as their meaning*. I could have expressed the same belief in a number of different ways, for example by saying *I've got a migraine*, or *my headache's come back*, or by clutching my head and saying *the usual problem again* in a long-suffering tone of voice. Since each of these utterances is expressed differently, they would correspond to different brain states. But we can concisely capture what they have in common by appealing to the level of their meaning: even though the brain states that produce them are different, they are united by the similarities of the meanings they convey. Just talking about brain states makes this elementary generalization impossible.

Especially at the current rudimentary stage of our knowledge of the brain, then, we have no choice but to continue to appeal to meanings in our explanations of language. Brain states are too complicated and too variable (both within and between individuals) to allow us to capture the straightforward generalizations we can capture using the intentional vocabulary of meaning. Meanings are the thread that guides us through the variety and confusion of brain states and input–output sequences; only by invoking meanings can we relate language to human behaviour and psychology in general. Understanding brain states will be important

for understanding language, but not at the expense of meaning. Studying brain states will tell us *how* the brain does what it does. Studying meaning as part of an intentional level study of human psychology and behaviour will tell us *what* it is doing and *why* it is doing it. It is thus a confusion of explanatory levels to claim that meaning can be reduced to brain state.

1.6.4 Meaning and use

An alternative to the three previous theories is the view that a word's meaning consists simply in the way it is used. This is the **use theory of meaning**, and it has been advanced, in different forms, by behaviourist psychologists such as Skinner (1957), and linguists such as Bloomfield (1933). (A rather different, non-behaviourist use theory was advanced by Wittgenstein 1953.) Behaviourist proponents of the use theory typically reject the very notion that words have hidden, unobservable properties called meanings: since meanings are inherently unobservable, it is, they would claim, unscientific to use them in explanations. (This argument would no longer be accepted by philosophers of science: scientific explanation *usually* involves unobservables.) Use theorists have claimed that the only objective, scientific way to explain language is to avoid postulating unobservable objects called meanings, and to attend only to what may actually be observed, the particular sequences of words and expressions that occur in actual examples of language use, and to describe the relation between these linguistic forms and the situations in which they are used. According to these investigators, the explanatory task of semantics is to provide not an abstract characterization of meanings, whether interpreted as concepts or denotations, but a causal, predictive account of the way a given language is actually used. In the words of Skinner (1957: 5), 'What happens when a man speaks or responds to speech is clearly a question about human behavior', and the only correct way to answer it is to proffer a precise account of what linguistic behaviour is likely to be produced in different situations.

Thus, for Bloomfield (1933: 139), the only *meaning* a linguistic form has is 'the situation in which the speaker utters it and the response which it calls forth in the hearer'. To take a particularly simple example, one of the 'meanings' of 'sorry' in English might be described as a situation where the speaker apologises; the hearer's typical response will be to treat the utterance as an apology and behave accordingly (e.g. by letting the incident drop, by not accusing the speaker of rudeness, by themselves saying sorry, etc.). We can describe this situation without having to make any reference to a 'meaning' of *sorry*: external analysis of the situation is all that is needed.

QUESTION Can the meanings of the following words be described in terms of situations? *Hi*, *please*, *you*, *apple*, *thanks*, *this*

The project of specifying the uses of linguistic units is not as remote as it might seem from the traditional semantic project of describing denotations or senses. Indeed, the traditional notion of meaning itself is

ultimately aimed at explaining language-use, since it is the meaning of individual linguistic expressions that is taken to explain the way they are used: words are used *in accordance with* their meanings. For proponents of the use theory of meaning, we should directly describe the actual situations themselves in which language is spoken or written, rather than doing this via the intermediary notion of meaning. When we have developed a full theory of the way in which speakers actually use language, then the goal of semantics will have been fulfilled.

The main objection against use theories of meaning is simply the mind-boggling variety of the situations in which linguistic forms may be used. As Bloomfield acknowledges, the number of different situations in which language is used is infinite. There are very few, if any, linguistic expressions which are automatically called up by a specifiable external situation. If the meaning of a linguistic form is the situation of the speaker's utterance and the hearer's response, there will be very few words for which a description like the one just given for *sorry* would even *seem* plausible. It would not seem to be a feasible project to specify the situations in which most of the words in the previous paragraph are used, since they are not highly context bound and can be used in practically any situation. Think of some of the possible situations in which the noun *way*, for example, might be used. To catalogue these, we would need to know the individual circumstances of a representative number of speaker/hearer pairs in whatever linguistic community we were investigating, including what was referred to by *way* on each occurrence of use, the situation which prompted the speaker to utter it, and the response given by the hearer. For even straightforward, unremarkable instances of *way* like *I don't know the way* or *which way is quicker?* this will already involve a huge variety of different specific situations. But if we add instances where *way* is used sarcastically, metaphorically, dishonestly, or simply by mistake, it will be clear that the use theory is massively complicated, and that the extraction of any regularities or generalizations about language use will be extremely complicated (Chomsky 1959 has classic objections against this kind of use theory of language).

The prospects for a use theory might be better if the focus changes from the individual word to higher-level linguistic units. It does seem to be the case that there are many phrases and sentences which have a more predictable relationship to their situations than the individual words of which they are composed. Thus, conversational routines like greetings, invitations, asking for the time, congratulating, wishing luck and many others involve highly stereotyped instances of language such as those in (44), which are to some extent predictable from the situations in which they occur.

(44) *how are you?*
do you have the time?
good luck
congratulations!
have a nice weekend.

Yet in spite of the perhaps greater possibilities at the phrase level, the problem for the use theory of meaning remains the enormous variety of sentences which make up any individual's linguistic behaviour. Even if there are some very stereotypical phrases which crop up more or less predictably in given situations, this does not detract from the huge number of phrases and sentences uttered by a language user which are novel. The use theory of meaning, in other words, seems to ignore the compositionality of language. It is *because* the meanings of sentences are built up out of the meanings of words that we can put words into different combinations to suit new communicative needs, including in situations which we have never previously encountered. The situations in which language is used are constantly changing, yet we do not mysteriously lose our ability to communicate. A theory of meaning must be able to explain how it is that we can use old words to convey *new* meanings which have never been previously conveyed, in situations in which we have never previously been placed.

QUESTION Do obsolete, old-fashioned or archaic words pose a problem for the use theory? If so, why? If not, why not? Do the conceptual and referential/denotational theories fare any better?

1.7 Meaning and explanation

We've now considered four proposals about the nature of meaning: meaning as reference/denotation, meaning as concepts, meaning as brain states and meaning as use. What conclusions can we draw? One particular conclusion concerns the status of the term 'meaning' itself. Even though the notion of a word's meaning can be used to facilitate many tasks on the level of practical language use (explanation of new words, translation from one language to another, prescriptive regulation of disputes over usage, etc.), and seems indispensable on the intentional level of explanation discussed in 1.6.3, we should consider the possibility that 'meaning' is essentially a *pretheoretical*, informal notion which will not have any precise equivalent in a detailed account of linguistic behaviour on the other two levels.

QUESTION What are some other everyday, pretheoretical notions about language which have to be abandoned for the purposes of 'scientific' linguistics?

Perhaps, then, we do not need to choose between the different theories of meaning discussed in the previous section. As suggested in 1.2, 'meaning' can be seen as a shorthand way of talking about a whole variety of separate phenomena which are all individually important in our talk about language, especially on the intentional level of explanation, but which do not necessarily correspond to any single entity that will be revealed by careful empirical study. The English language category 'meaning', in other words, which in any case only has approximate equivalents in other languages, might have no precise role in a full understanding of language.

By contrast, the various aspects of meaning that we have distinguished in this chapter – reference, conceptual content, connotation, and so on – are all factors for which linguistic semantics *does* owe a principled explanation, and for which it should try to find theoretical analogues. There is no single way of breaking the definitional circle: 'meaning' is many different things, none of which should be ruled out as irrelevant to the eventual explanation of language.

In this context, we can specify an important condition that any principled theory of language must meet. This condition is linked to the idea that the ultimate goal of research into language must be to contribute to the *causal explanation* of people's utterances. To achieve a thorough understanding of our linguistic ability, we will eventually need to be able to specify the detailed causal mechanisms which lead up to the production of utterances by speakers in real time. In order to achieve this, we will need a precise account of what the various phonological, semantic, morphosyntactic and semantic properties of different linguistic forms are, and of the ways in which these properties are combined in actual discourse sequences in different contexts of use. This goal is exactly the same as the one aimed at in other sciences: chemistry, for example, specifies the various properties of different molecules, in virtue of which they enter into sequences of causal interaction with each other, and embryology aims at understanding the properties of fertilized cells, in virtue of which a step-by-step understanding of their development into full organisms can be achieved. In the case of linguistics, the detailed nitty-gritty of a causal account is a long way off. What is more, the fine detail of an account of linguistic behaviour on the implementational level will have to be provided by neurolinguists and other brain scientists who will be able to isolate the physiological underpinnings of linguistic phenomena. The semanticist's role is an earlier one, which consists in isolating the important properties of the linguistic system, on the intentional and perhaps algorithmic levels, for which these experimental scientists will need to find the physical mechanisms.

The distant goal of a causal account of language behaviour, however, suggests a possible role for the notion of meaning in semantics. From this perspective, we can suggest that to talk about a word's meaning is a shorthand way of talking about *whatever property of a word could enter into causal explanations of its use*. In our ordinary talk about language, one of the main functions of the category of meaning is to explain word use: we use the words we use because of the meanings they have. But in order to go beyond this pretheoretical level and explain a word's use in a rigorous way, which might be ideally compatible with a causal account of language, a word's meaning may include many different explanatory properties and necessitate consideration of referents *and* concepts *and* situations of use.

As a result, we do not need any single, categorical answer to the question of whether meaning is denotation *or* concepts *or* uses. To phrase the question as a set of exclusive choices like this is counterproductive, since it may well turn out that *all* of these categories will need to be invoked in order to explain the use of different words. Thus, as we noted at (44), there is a subclass of words and phrases in any language whose use seems particularly

closely linked to certain recurrent and specifiable situations; the most obvious way of explaining the use of these words is to associate them with the particular contexts and situations in which they occur, and the use theory of meaning will be the most relevant. Other words, however, seem best explained by the particular conceptual associations they call up; for these, attention to the link between words and concepts will be the most relevant. If I say, for example, *The holidays were a nightmare*, then the words *holidays* and *nightmare* call up a whole variety of specific connotations and associations (see the question below) for which the conceptual theory of meaning will be most appropriate. In still other cases, such as proper names and 'deictics' like *here*, it seems to be a word's referent which is the most important factor in accounting for the word's use on a given occasion: if I say *that man just fell over*, the 'meaning' of *that man* is best described as the actual person to whom I am referring. This is not to say that concepts are irrelevant for expressions like *that man* or for words like those in (44), or that referents and denotations are irrelevant for words like *holiday* or *nightmare*. In most cases, indeed, we will need to attend to all three aspects of a word's 'meaning', in considering how its relations with referents/denotations, associated concepts and uses *mutually* combine to account for its presence in a particular linguistic context. It is just to say that in all these cases attention to the explanatory purpose of talk about meaning will direct us towards whichever conception of meaning seems to provide the best explanation of the particular semantic phenomenon at hand.

QUESTION Describe the concepts HOLIDAY and NIGHTMARE in as much detail as possible. How much of this detail is relevant to explaining linguistic behaviour?

Summary

The meaningfulness of language is an instance of the meaningfulness of behaviour

The meaningfulness of language can be seen as just one instance of the meaningfulness of human behaviour and communication in general, and is one of the systems of structured meaningfulness studied in semiotics.

'Meaning' is a very vague term

'Meaning' is a very vague term: in English it refers to a variety of different relations between the world, language and speakers. Most languages do not have precise equivalents for the English term 'meaning', and some use a very different stock of lexical resources to talk about meaning-like phenomena.

The semiotic triangle

For the purposes of linguistics, we can isolate three particularly important factors relevant to the study of meaning: the **psychology** of

speakers, which creates and interprets language, the **referent** of the language expression as projected by the language user's psychology, and the **linguistic expression** itself: these three points constitute the **semiotic triangle**.

Lexemes

In providing a semantic description of a language, we do not need to treat all the variant morphological forms of a single word separately. Instead, we describe the meanings of a language's **lexemes**, or the abstract units which unite all the morphological variants of a single word.

Sense, reference, denotation and connotation

There are several different aspects of the meaning of a lexeme: its **referent** on any one occasion of use, its **denotation**, which is the set of all its referents, and its **sense**, or the abstract, general meaning which can be translated from one language to another, paraphrased, or defined in a dictionary. **Connotation** names those aspects of meaning which do not affect a word's sense, reference or denotation, but which have to do with secondary factors such as its emotional force, its level of formality, its character as a euphemism, etc.

Compositionality

Meaning is often **compositional**, which means that the meanings of sentences are made up, or **composed**, of the meanings of their constituent lexemes.

Sentence and utterance meaning

Sentence meaning is the compositional meaning of the sentence as constructed out of the meanings of its individual component lexemes. **Utterance meaning** is the meaning which the words have on a particular occasion of use in the particular context in which they occur. **Semantics** studies sentence meaning, whereas **pragmatics** studies utterance meaning and other aspects of language use.

Object language and metalanguage

In analysing meaning we distinguish the **object language**, or the language *whose* meanings are being described, from the **metalanguage**, the language *in which* we describe these meanings.

Explanations of meaning in terms of meanings are circular

When we propose a definition in a metalanguage as an analysis of the meaning of an object language term, the more basic questions, 'what is meaning?' and 'what is it to understand a meaning?' are left unanswered. All definitions of meaning in language, therefore, are ultimately **circular** because they use one kind of meaning to explain another.

Four ways of breaking the circle

There are four important answers to the question 'what is meaning?': the **referential/denotational theory** of meaning, the **conceptual theory** of meaning, the **brain states theory** and the **use theory**. We do not have to categorically choose between these theories. Instead, recognizing that the notion of meaning in linguistics is a way of talking about *the factors which explain language use,* we can see referents, concepts, brain states and uses as all relevant to this task.

Further reading

Saussure (1983) is essential reading for semantics, as for linguistics generally. For useful introductions to semiotics, see Sebeok (1994), Cobley (2001) and Hawkes (1983). Lyons (1977: Chapter 7) provides a thorough introduction to the concepts of sense and reference; see also Chapter 3 of this book and the references mentioned there. Levinson (1983) and Mey (2001) are standard introductions to utterance meaning and pragmatics. Martin (1987), Frawley (1992) and Chapter 2 of Allan (1986) are good introductions to different theories of meaning. On the role of concepts in semantics see Jackendoff (1983) and (1989) and on concepts more generally, the opening chapters of Prinz (2002). Cummins (1989) is an introduction to meaning and mental representation, and Murphy (2002) is a compendium of psychological research on concepts, including their relation to word meaning. Lakoff (1987) explores a specific conceptual theory of meaning. Lyons (1977: Chapter 5) is a detailed account of the use theory of meaning. Jung-Beeman (2005) gives a glimpse into research on meaning in cognitive neuroscience. On the history of modern European and American semantics, see Gordon (1982). For information on non-European semantic traditions and a discussion of the Greek origins of Western semantics, see van Bekkum *et al.* (1997). Ullmann (1972) and Ogden and Richards (1949) are classic works in the history of semantics which still have many insights. On the contrast between theoretical and pretheoretical perspectives in linguistics, see Chomsky (2000).

Exercises

Questions for discussion

1. In Section 1.1 we discussed the relation between meaning, communication and significance. Consider the cases of pure, wordless music and 'nonsense' language. Can either of these be said to be meaningful? If so, how is this meaningfulness different from that of language? Would you consider it as communication? If so, what is communicated? If not, why not?
2. In 1.2 we considered the words available for the representation of meaning-phenomena in English, French, Warlpiri and Chinese. Choose a language you know and describe what words are available to talk about meaning, and their similarities and differences with the languages discussed.
3. In ancient philosophy, the study of the meanings of words was not usually recognized as a distinct subject. Instead, language and meaning were mainly discussed for what they revealed about the nature of the world, logic and our ideas. What do you think the most important links are between the study of linguistic semantics and other branches of enquiry?
4. We saw in 1.6.1 that some linguistic expressions have a sense but do not have a reference/denotation. Do you think there could be any linguistic expressions with reference/denotation but no senses? If so, what are they? If not, why not?

5. For each word in the following sentences, assume a particular occasion of utterance, and try to specify a sense, referent and denotation. Is this possible for all the words? Could it make sense to talk about the referents of words which are not nouns? Do any words have particular connotations?

   (a) The standard incubation period of Mad Cow Disease is between three and five years.
   (b) You are squabbling about the question whether the buttons of the National Guard should be white or yellow.
   (c) Brazilian officials, like those from India to China, describe their steadily expanding space effort as a commercial and strategic necessity, as well as a matter of national prestige.
   (d) You gave that one to me then.

6. The phrase *Australian passport* is clearly compositional, since its meaning consists of the meaning of *passport* and the meaning of *Australian*: an *Australian passport* is a passport that is Australian. But consider the phrase *a false passport.* A false passport is not a passport that is false, since a false passport is not a (real) passport at all. There is thus a way in which the meaning of the phrase *false passport* does not contain the meaning of the word *passport*. Is this a problem or not for the idea of compositionality? Are the following phrases also problematic? If so, state why. If not, why not?

   This is ***fake caviar.***
   My ***old baby-sitter*** was only eleven.
   A flea is less than a ***millimetre high***.
   I am going ***partly bald***.

7. Like other branches of linguistics, semantics is a *descriptive*, not a *prescriptive* enterprise, and aims to describe the meanings of words as they are actually used by speakers, and not as they 'should' be used. Give examples of, and describe the meanings of the following words, and comment on any discrepancies between this description and a prescriptive view of their meaning: *disinterested*, *infer*, *fulsome*, *inflammable*, *champagne*, *monkey*, *insane*, *golden.*
8. Consider the following quotation:

   We must not allow our words to change their meanings, but must make sure that we use them in their correct senses. For if we are careless with meanings, we will lose them, and there will be many ideas which we will no longer be able to express. For 'disinterested' does not mean the same as 'uninterested', 'fulsome' does not mean the same as 'full', 'infer' does not mean the same as 'imply'. If we lose these differences of meaning, we will lose the differences in the concepts they express.

   Do you agree with these statements? What assumptions about language do they contain?
9. Some dictionaries use pictures in order to escape the problem of circular definitions. What are the advantages and limitations of this strategy? Consider how easily the meaning of the following words could be conveyed pictorially: *oak*, *to punch*, *black*, *happy*, *microscope*, *water*, *underneath*, *arch*, *machine*, *sensitivity*, *internet*, *thin*, *popular*, *to sleep*, *horrendous*.

10. The Internet contains a number of automatic translation programmes. What problems do you think are posed by the project of translating meanings from one language to another automatically? Why are human beings so much better at translation than computers? Consider the role of context, background knowledge, intuition, memory and any others which you think are relevant.
11. We often talk about seeing or putting meanings into things which do not have them. Astrologers, for example, see 'meanings' in the stars. Other people claim to understand meanings in tea leaves. How are these 'meanings' different from the ones communicated in language?
12. Words are not the only linguistic units that communicate meaning. Discuss whether each of the following categories can communicate meaning, and, if so, what sorts: intonation, speech volume, speech speed, length of sentence, choice of language, and (for type-written language) choice of typeface.
13. Does it make sense to speak of a lexeme's reference?
14. Consider the following words: *squabble*, *fight*, *argue*, *bicker*, *dispute*, *disagree*, *debate*, *contend*, *spat*. Describe how they differ in sense and connotation.
15. Review the examples in (40). Now try to devise alternative descriptions of the denotations of the following nouns: *hand*, *baseball*, *breakfast*, *nine*, *red*, *stranger*, *person*, *heart disease*.
16. You are sent out to learn, and write a dictionary and grammar of a previously unrecorded language. While in the field, you notice that whenever a plane passes overhead the speakers look up and utter the word *paabo*. How many different possibilities can you think of for the meaning of this word? What problems can you imagine in trying to work out which is the right one? Are there any general consequences for the study of meaning?