Preliminaries

part I

chapter 1

Semantics in Linguistics

1.1 Introduction

Semantics is the study of meaning communicated through language. This book is an introduction to the theory and practice of semantics in modern linguistics. Although this is not an introduction to any single theory, we begin with a basic assumption: that a person's linguistic abilities are based on knowledge that they have. It is this knowledge that we are seeking to investigate. One of the insights of modern linguistics is that speakers of a language have different types of linguistic knowledge, including how to pronounce words, how to construct sentences, and about the meaning of individual words and sentences. To reflect this, linguistic description has different **levels of analysis**. So **phonology** is the study of what sounds a language has and how these sounds combine to form words; **syntax** is the study of how words can be combined into sentences; and **semantics** is the study of the meanings of words and sentences.

The division into levels of analysis seems to make sense intuitively: if you are learning a foreign language you might learn a word from a book, know what it means but not know how to pronounce it. Or you might hear a word, pronounce it perfectly but not know what it means. Then again, you might know the pronunciation and meaning of, say, a noun, but not know

how its plural is formed or what its genitive case looks like. In this sense knowing a word unites different kinds of knowledge, and this is just as true of your knowledge of how to construct phrases and sentences.

Since linguistic description is an attempt to reflect a speaker's knowledge, the semanticist is committed to describing semantic knowledge. This knowledge allows English speakers to know, for example, that both the following sentences describe the same situation:

1.1   In the spine, the thoracic vertebrae are above the lumbar vertebrae.

1.2   In the spine, the lumbar vertebrae are below the thoracic vertebrae.

that 1.3 and 1.4 below **contradict** each other:

1.3   Addis Ababa is the capital of Ethiopia.

1.4   Addis Ababa is not the capital of Ethiopia.

that 1.5 below has several possible meanings, i.e. is **ambiguous**:

1.5   She gave her the slip.

that 1.6 below **entails** 1.7:

1.6   Henry murdered his bank manager.

1.7   Henry's bank manager is dead.

We will look at these types of semantic knowledge in more detail a little later on; for now we can take **entailment** to mean a relationship between sentences so that if a sentence *A* entails a sentence *B*, then if we know *A* we automatically know *B*. Or alternatively, it should be impossible, at the same time, to assert *A* and deny *B*. Knowing the effect of inserting the word *not*, or about the relationships between *above* and *below*, and *murder* and *dead*, are aspects of an English speaker's semantic knowledge, and thus should be part of a semantic description of English.

As our original definition of semantics suggests, it is a very broad field of inquiry, and we find scholars writing on very different topics and using quite different methods, though sharing the general aim of describing semantic knowledge. As a result semantics is the most diverse field within linguistics. In addition, semanticists have to have at least a nodding acquaintance with other disciplines, like philosophy and psychology, which also investigate the creation and transmission of meaning. Some of the questions raised in these neighbouring disciplines have important effects on the way linguists do semantics. In chapter 2 we discuss some of these

questions, but we begin in this chapter by looking at the basic tasks involved in establishing semantics as a branch of linguistics.

1.2   Semantics and Semiotics

So we see our basic task in semantics as showing how people communicate meanings with pieces of language. Note, though, that this is only part of a larger enterprise of investigating how people understand meaning. Linguistic meaning is a special subset of the more general human ability to use signs, as we can see from the examples below:

1.8   Those vultures mean there's a dead animal up ahead.

1.9   His high temperature may mean he has a virus.

1.10   The red flag means it's dangerous to swim.

1.11   Those stripes on his uniform mean that he is a sergeant.

The verb *mean* is being put to several uses here, including inferences based on cause and effect, and on knowledge about the arbitrary symbols used in public signs. These uses reflect the all-pervasive human habit of identifying and creating signs: of making one thing stand for another. This process of creating and interpreting symbols, sometimes called **signification**, is far wider than language. Scholars like Ferdinand de Saussure (1974) have stressed that the study of linguistic meaning is a part of this general study of the use of sign systems, and this general study is called **semiotics**.[1] Semioticians investigate the types of relationship that may hold between a sign and the object it represents, or, in de Saussure's terminology, between a **signifier** and its **signified**. One basic distinction, due to C. S. Peirce, is between **icon**, **index** and **symbol**. An icon is where there is a similarity between a sign and what it represents, as for example between a portrait and its real-life subject, or a diagram of an engine and the real engine. An index is where the sign is closely associated with its signified, often in a causal relationship; thus smoke is an index of fire. Finally, a symbol is where there is only a conventional link between the sign and its signified, as in the use of insignia to denote military ranks, or perhaps the way that mourning is symbolized by the wearing of black clothes in some cultures, and white clothes in others. In this classification, words would seem to be examples of verbal symbols.[2]

In our discussion of semantics we will leave this more comprehensive level of investigation and concentrate on linguistic meaning. The historical development between language and other symbolic systems is an open question: what seems clear is that language represents man's most sophisticated use of signs.

1.3 Three Challenges in Doing Semantics

Analysing a speaker's semantic knowledge is an exciting and challenging task, as we hope to show in this book. We can get some idea of how challenging by adopting a simple but intuitively attractive theory of semantics which we can call the **definitions theory**. This theory would simply state that to give the meaning of linguistic expressions we should establish definitions of the meanings of words. We could then assume that when a speaker combines words to form sentences according to the grammatical rules of her[3] language, the word definitions are combined to form phrase and then sentence definitions, giving us the meanings of sentences. Let us investigate putting this approach into practice.

As soon as we begin our task of attaching definitions to words, we will be faced with a number of challenges. Three in particular prove very tricky for our theory. The first is the problem of **circularity**. How can we state the meaning of a word, except in other words, either in the same or a different language? This is a problem that faces dictionary writers: if you look up a word like *ferret* in a monolingual English dictionary, you might find a definition like 'Domesticated albino variety of the polecat, *Mustela putorius*, bred for hunting rabbits, rats, etc.' To understand this, you have to understand the words in the definition. According to our aims for semantics, we have to describe the meanings of these words too, beginning with *domesticated*. The definition for this might be 'of animals, tame, living with human beings'. Since this definition is also in words, we have to give the meaning, for example, of *tame*. And so on. If the definitions of word meaning are given in words, the process might never end. The question is: can we ever step outside language in order to describe it, or are we forever involved in circular definitions?

A second problem we will meet is how to make sure that our definitions of a word's meaning are exact. If we ask where the meanings of words exist, the answer must be: in the minds of native speakers of the language. Thus meaning is a kind of knowledge. This raises several questions; for example: is there a difference between this kind of knowledge and other kinds of knowledge that people have? In particular: can we make a distinction between **linguistic knowledge** (about the meaning of words) and **encyclopaedic knowledge** (about the way the world is)? For example, if I believe that a whale is a fish, and you believe that it is a mammal, do our words have different meanings when we both use the noun *whale*? Presumably you still understand me when I say *I dreamt that I was swallowed by a whale*.

There is another aspect to this problem: what should we do if we find that speakers of a language differ in their understanding of what a word means? Whose knowledge should we pick as our 'meaning'? We might avoid the decision by picking just one speaker and limiting our semantic description to an **idiolect**, the technical term for an individual's language. Another strategy to resolve differences might be to identify experts and use their

knowledge, but as we shall see, moving away from ordinary speakers to use a scientific definition for words has the danger of making semantics equivalent to all of science. It also ignores the fact that most of us seem to understand each other talking about, say, animals, without any training in zoology. This is a point we will come back to in chapter 2.

A third type of challenge facing us comes from looking at what particular utterances mean in context. For example: if someone says to you *Marvellous weather you have here in Ireland*, you might interpret it differently on a cloudless sunny day than when the rain is pouring down. Similarly *He's dying* might mean one thing when said of a terminally ill patient, and another as a comment watching a stand-up comedian failing to get laughs. Or again: *It's getting late* if said to a friend at a party might be used to mean *Let's leave*. The problem here is that if features of context are part of an utterance's meaning then how can we include them in our definitions? For a start, the number of possible situations, and therefore of interpretations, is enormous if not infinite. It doesn't seem likely that we could fit all the relevant information into our definitions.

These three issues: circularity; the question of whether linguistic knowledge is different from general knowledge; and the problem of the contribution of context to meaning, show that our definitions theory is too simple to do the job we want. Semantic analysis must be more complicated than attaching definitions to linguistic expressions. As we shall see in the rest of this book, semanticists have proposed a number of strategies for improving on this initial position. In the next section we discuss some initial ideas that will enable us to follow these strategies.

1.4 Meeting the Challenges

In most current linguistic theories, semantic analysis is as important a part of the linguist's job as, say, phonological analysis. Theories differ on details of the relationship between semantics and other levels of analysis like syntax and morphology, but all seem to agree that linguistic analysis is incomplete without semantics. We need, it seems, to establish a semantic component in our theories. We have to ask: how can we meet the three challenges outlined in the last section? Clearly we have to replace a simple theory of definitions with a theory that successfully solves these problems.

One of the aims of this book is to show how various theories have sought to provide solutions to these problems and we will return to them in detail over subsequent chapters. For now we will simply mention possible strategies which we will see fleshed out later. To cope with the problem of circularity, one solution is to design a semantic **metalanguage** with which to describe the semantic units and rules of all languages. We use metalanguage here with its usual meaning in linguistics: the tool of description. So in a grammar of Arabic written in French, Arabic is the *object language* and

French the *metalanguage*. An ideal metalanguage would be neutral with respect to any natural languages, i.e. would not be unconsciously biased towards English, French, etc. Moreover it should satisfy scientific criteria of clarity, economy, consistency, etc. We will see various proposals for such a metalanguage, for example to represent word meanings and the semantic relations between words, in chapters 9 and 10. We will also meet claims that such a metalanguage is unattainable and that the best policy is to use ordinary language to describe meaning.

For some linguists, though, translation into even a perfect metalanguage would not be a satisfactory semantic description. Such a line of reasoning goes like this: if words are symbols, they have to relate to something; otherwise what are they symbols of? In this view, to give the semantics of words we have to ground them in something non-linguistic. In chapter 2 we will review the debate about whether the things that words signify are real objects in the world or thoughts.

Setting up a metalanguage might help too with the problem of relating semantic and encyclopaedic knowledge, since designing meaning representations, for example for words, involves arguing about which elements of knowledge should be included. To return to our earlier example of *whale*: we assume that English speakers can use this word because they know what it means. The knowledge a speaker has of the meaning of words is often compared to a mental **lexicon** or dictionary. Yet if we open a real dictionary at the entry for *whale*, the definition is likely to begin 'large marine mammal . . .'. To rephrase our earlier question: does it follow that someone who doesn't know that whales are mammals fails to understand the meaning of the word *whale*? What if the speaker knows that it is a large animal that lives in the sea, but is hazy after that? The real issue is the amount of knowledge that it is necessary to know in order to use a word. We shall see aspects of this debate, which is really part of the general psychological debate about the representation of concepts and categories, in chapters 2, 3 and 7.

In tackling the third problem, of context, one traditional solution has been to assume a split in an expression's meaning between the local contextual effects and a context-free element of meaning, which we might call **conventional** or **literal** meaning. We could perhaps try to limit our definitions to the literal part of meaning and deal with contextual features separately. As we shall see in chapter 3, though, it turns out to be no easy task to isolate the meaning of a word from any possible context. We discuss some aspects of this idea of literal meaning in 1.6.3 below. The other side of such an approach is to investigate the role of contextual information in communication, and try to establish theories of how speakers amalgamate knowledge of context with linguistic knowledge. As we shall see in chapter 7, it seems that speakers and hearers cooperate in using various types of contextual information. Investigating this leads us to a view of the listener's role which is quite different from the simple, but common, analogy of decoding a coded message. We shall see that listeners have a very active role, using what has been said, together with background knowledge, to make inferences

about what the speaker meant. The study of these processes and the role in them of context is often assigned to a special area of study called **pragmatics**. We discuss the relationship between semantics and pragmatics in 1.6.4 below. We shall see instances of the role of context in meaning throughout this book and this will give us the opportunity to review the division of labour between semantics and this newer field of pragmatics.[4]

Each of these strategies will be investigated in later chapters of this book: the creation of semantic metalanguages, the modelling of conceptual knowledge, the theory of literal language, and factoring out context into pragmatics. Meanwhile in the next section we look at how semantics might fit into a model of language.

1.5　Semantics in a Model of Grammar

1.5.1　Introduction

As has been suggested already, for many linguists the aim of doing semantics is to set up a component of the grammar which will parallel other components like syntax or phonology. Linguists like to draw flowchart-style diagrams of grammatical models, and in many of them there is a box labelled 'semantics', as in figure 1.1:

Figure 1.1　Components of grammar

Before we go on, it might be worthwhile to consider whether it is justified to view semantics as a component equal and parallel to, say, syntax.

We saw earlier that linguists identify different levels of analysis. Another way of describing this is to say that linguistic knowledge forms distinct **modules**, or is **modularized**. As a result, many linguistic theories are themselves modularized, having something like our boxes in figure 1.1. Our question, though, remains: what kind of module is semantics? The answer varies from theory to theory. The real problem is, of course, that units at all linguistic levels serve as part of the general enterprise: to communicate meaning. This means that in at least one sense, meaning is a product of all linguistic levels. Changing one phoneme for another, one verb ending for another, or one word order for another will produce differences of meaning. This view leads some writers to believe that meaning cannot be identified as a separate level, autonomous from the study of other levels of grammar. A strong version of this view is associated with the theory known as **cognitive grammar**, advocated by linguists such as Ronald Langacker (e.g. Langacker 1987);[5] see, for example, this claim from a recent collection of articles:

1.12   the various autonomy theses and dichotomies proposed in the lin-
        guistic literature have to be abandoned: a strict separation of syn-
        tax, morphology and lexicon is untenable; furthermore it is
        impossible to separate linguistic knowledge from extra-linguistic
        knowledge.   (Rudzka-Ostyn 1993: 2)

As we shall see in the course of this book, however, many other linguists
do see some utility in maintaining both types of distinction referred to
above: between linguistic and non-linguistic knowledge; and within linguistic
knowledge, identifying distinct modules for knowledge about pronunciation,
grammar and meaning.

1.5.2   Word meaning and sentence meaning

If an independent component of semantics is identified, one central issue is
the relationship between word meaning and sentence meaning. Knowing a
language, especially one's native language, involves knowing thousands of
words. As mentioned earlier, we can call the mental store of these words a
lexicon, making an overt parallel with the lists of words and meanings
published as dictionaries. We can imagine the mental lexicon as a large but
finite body of knowledge, part of which must be semantic. This lexicon is
not completely static because we are continually learning and forgetting
words. It is clear, though, that at any one time we hold a large amount of
semantic knowledge in memory.

   Phrases and sentences also have meaning, of course, but an important dif-
ference between word meaning on the one hand, and phrase and sentence
meaning on the other, concerns **productivity**. It is always possible to create
new words, but this is a relatively infrequent occurrence. On the other hand,
speakers regularly create sentences that they have never used or heard before,
confident that their audience will understand them. Noam Chomsky in par-
ticular has commented on the creativity of sentence formation (for example
Chomsky 1965: 7–9). It is one of generative grammar's most important
insights that a relatively small number of combinatory rules may allow speakers
to use a finite set of words to create a very large, perhaps infinite, number
of sentences. To allow this the rules for sentence formation must be **recursive**,
allowing repetitive embedding or coordination of syntactic categories. To
give a simple example, a compositional rule like 1.13 below, where elements
in parentheses are optional and the asterisk means the optional group is
repeatable, will allow potentially limitless expansions of S, as in 1.14:

1.13   S → [$_S$ S (and S)*]

1.14   a.   [$_S$ S and S]
        b.   [$_S$ S and S and S]]
        c.   [$_S$ S and S and S and S] etc.

The idea is that you can always add another clause to a sentence, or, as 1.15 and 1.16 below show, another nominal within a nominal:

1.15   NP → [$_{NP}$ NP (and NP)*]

1.16   a.   I bought [$_{NP}$ a book]
        b.   I bought [$_{NP}$ [$_{NP}$ a book] and [$_{NP}$ a magazine]]
        c.   I bought [$_{NP}$ [$_{NP}$ a book] and [$_{NP}$ a magazine] and [$_{NP}$ some pens]] etc.

See Lyons (1968: 221–2) for discussion of such recursive rules in syntax.

   This insight has implications for semantic description. Clearly, if a speaker can make up novel sentences and these sentences are understood, then they obey the semantic rules of the language. So the meanings of sentences cannot be listed in a lexicon like the meanings of words: they must be created by rules of combination too. Semanticists often describe this by saying that sentence meaning is **compositional**. This term means that the meaning of an expression is determined by the meaning of its component parts and the way in which they are combined.

   This brings us back to our question of levels. We see that meaning is in two places, so to speak, in a model of grammar: a more stable body of word meanings in the lexicon, and the limitless composed meanings of sentences. How can we connect semantic information in the lexicon with the compositional meaning of sentences? It seems reasonable to conclude that semantic rules have to be compositional too and in some sense 'in step' with grammatical rules. The relationship is portrayed differently in different theories of language. In the evolving forms of Noam Chomsky's generative grammar (e.g. Chomsky 1965, 1988) syntactic rules operate independently of semantic rules but the two types are brought together at a level of Logical Form. In many other theories semantic rules and grammatical rules are inextricably bound together, so each combination of words in a language has to permissible under both. Such an approach is typical of functional approaches like Halliday's Functional Grammar (1994), and Role and Reference Grammar (Van Valin and LaPolla 1997), as well as variants of generative grammar like Head-Driven Phrase Structure Grammar (Sag and Wasow 1999).[6]

1.6   Some Important Assumptions

At this point we can introduce some basic ideas that are assumed in many semantic theories and that will come in useful in our subsequent discussion. In most cases the descriptions of these ideas will be simple and a little on the vague side: we will try to firm them up in subsequent chapters.

Figure 1.2 Reference and sense in the vocabulary

LINGUISTIC VALUE
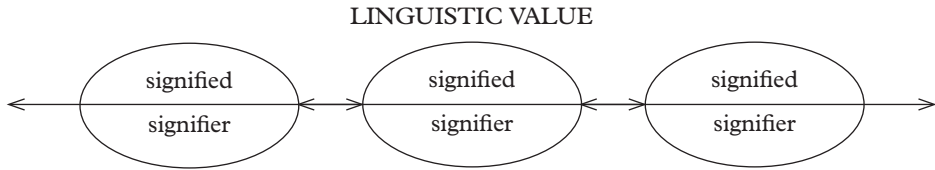
1.6.1 Reference and sense

One important point made by the linguist Ferdinand de Saussure (1974), whose ideas have been so influential in the development of modern linguistics, is that the meaning of linguistic expressions derives from two sources: the language they are part of and the world they describe. Words stand in a relationship to the world, or our mental classification of it: they allow us to identify parts of the world, and make statements about them. Thus if a speaker says _He saw Paul_ or _She bought a dog_, the underlined nominals identify, pick out, or **refer** to specific entities in the world. However words also derive their value from their position within the language system. The relationship by which language hooks onto the world is usually called **reference.** The semantic links between elements within the vocabulary system are an aspect of their **sense**,[7] or meaning.

Ferdinand de Saussure (1974: 115) used the diagram in figure 1.2 to show this patterning. Each oval is a word, having its own capacity for reference, but each is also linked to other words in the same language, like a cell in a network. His discussion of this point is excellent and we cannot really do it justice here, except to recommend the reader to the original. His well-known examples include a comparison of English _sheep_ and French _mouton_. In some cases they can be used to refer in a similar way but their meaning differs because they are in different systems and therefore have different ranges: in English there is an extra term _mutton_, used for meat, while the French word can be used for both the animal and the meat. Thus, the meaning of a word derives both from what it can be used to refer to and from the way its semantic scope is defined by related words. So the meaning of _chair_ in English is partly defined by the existence of other words like _stool_. Similarly, the scope of _red_ is defined by the other terms in the colour system: _brown_, _orange_, _yellow_, etc. The same point can be made of grammatical systems: de Saussure pointed out that plural doesn't 'mean' the same in French, where it is opposed to singular, as it does in Sanskrit or Arabic, languages which, in addition to singular, have **dual** forms, for exactly two entities. In the French system, plural is 'two or more'; in the other systems, 'three or more'.

1.6.2 Utterances, sentences and propositions

These three terms are used to describe different levels of language. The most concrete is **utterance**: an utterance is created by speaking (or writing)

a piece of language. If I say *Ontogeny recapitulates phylogeny*, this is one utterance. If another person in the same room also says *Ontogeny recapitulates phylogeny*, then we would be dealing with two utterances.

Sentences, on the other hand, are abstract grammatical elements obtained from utterances. Sentences are abstract because if a third and fourth person in the room also say *Ontogeny recapitulates phylogeny* with the same intonation, we will want to say that we have met four utterances of the same sentence. In other words, sentences are abstracted, or generalized, from actual language use. One example of this abstraction is direct quotation. If someone reports *He said 'Ontogeny recapitulates phylogeny'*, she is unlikely to mimic the original speaker exactly. Usually the reporter will use her normal voice and thus filter out certain types of information: the difference in pitch levels between men, women and children; perhaps some accent differences due to regional or social variation; and certainly those phonetic details which identify individual speakers. Speakers seem to recognize that at the level of the sentence these kinds of information are not important, and so discard them. So we can look at sentences from the point of view of the speaker, where they are abstract elements to be made real by uttering them; or from the hearer's point of view, where they are abstract elements reached by filtering out certain kinds of information from utterances.

One further step of abstraction is possible for special purposes: to identify **propositions**. In trying to establish rules of valid deduction, logicians discovered that certain elements of grammatical information in sentences were irrelevant; for example, the difference between active and passive sentences:

1.17   Caesar invaded Gaul.

1.18   Gaul was invaded by Caesar.

From a logician's perspective, these sentences are equivalent, for whenever 1.17 is true, so is 1.18. Thus the grammatical differences between them will never be significant in a chain of reasoning and can be ignored. Other irrelevant information (for these purposes) includes what we will in chapter 7 call **information structure**, i.e. the difference between the following sentences:

1.19   It was Gaul that Caesar invaded.

1.20   It was Caesar that invaded Gaul.

1.21   What Caesar invaded was Gaul.

1.22   The one who invaded Gaul was Caesar.

These sentences seem to share a description of the same state of affairs. Once again, if one is true all are true, and if one is false then all are false.

To capture this fact, logicians identify a common proposition. Such a proposition can be represented in various special ways to avoid confusion with the various sentences which represent it, e.g. by using capitals:

1.23   CAESAR INVADED GAUL.

Thus the proposition underlying the sentence *The war ended* might be written:

1.24   THE WAR ENDED.

Logicians commonly use formulae for propositions in which the verb is viewed as a function, and its subject and any objects as arguments of the function. Such formulae often delete verb endings, articles and other grammatical elements, so that corresponding to 1.23 and 1.24 we would get 1.25 and 1.26 below:

1.25   invade (Caesar, Gaul)

1.26   end (war)

Some semanticists have borrowed from logicians both this notion of proposition and the use of logical formulae. We will see various applications of such formulae in later chapters.[8] As we shall see, some linguists employ this notion of proposition in their semantic analysis, often to identify a description of an event or situation which might be a shared element in different sentences. So, for example, the statement *Joan made the sorbet*, the question *Did Joan make the sorbet?*, and the command: *Joan, make the sorbet!* might be seen to share a propositional element: JOAN MAKE THE SORBET. In this view, these different sentences allow the speaker to do different things with the same proposition: to assert it as a past event; to question it; or to request someone to bring it about.

Propositions, then, can be a way of capturing part of the meaning of sentences. They are more abstract than sentences because, as we saw in examples 1.17–22 above, the same proposition can be represented by several different statements. Moreover in non-statements like questions, orders, etc. they cannot be the complete meaning since such sentences include an indication of the speaker's attitude to the proposition. We will come back to the linguistic marking of such attitudes in chapter 8.

To sum up: **utterances** are real pieces of speech. By filtering out certain types of (especially phonetic) information we can get to abstract grammatical elements, **sentences**. By going on to filter out certain types of grammatical information, we can get to **propositions**, which are descriptions of states of affairs and which some writers see as a basic element of sentence meaning. We will get some idea of the different uses to which these terms are put in the remainder of this book.[9]

1.6.3 Literal and non-literal meaning

This distinction is assumed in many semantics texts but attempting to define it soon leads us into some difficult and theory-laden decisions. The basic distinction seems a common-sense one: distinguishing between instances where the speaker speaks in a neutral, factually accurate way, and instances where the speaker deliberately describes something in untrue or impossible terms in order to achieve special effects. Thus if one afternoon you were feeling the effects of missing lunch, you might speak literally as in 1.27, or non-literally as in 1.28–30:

1.27   I'm hungry.

1.28   I'm starving.

1.29   I could eat a horse.

1.30   My stomach thinks my throat's cut.

Non-literal uses of language are traditionally called **figurative** and are described by a host of rhetorical terms including **metaphor**, **irony**, **metonymy**, **synecdoche**, **hyperbole** and **litotes**. We will meet examples of these terms later on. On closer examination, though, it proves difficult to draw a firm line between literal and non-literal uses of language. For one thing, one of the ways languages change over time is by speakers shifting the meanings of words to fit new conditions. One such shift is by metaphorical extension, where some new idea is depicted in terms of something more familiar. For a while the new expression's metaphorical nature remains clear, as for example in the expressions *glass ceiling* for promotional barriers to women, or *surfing* the internet. Slightly older coinings might include *mouse* for the computer keyboard extension, or expressions like *toy boy*, or *junk bonds*. After a while such expressions become fossilized and their metaphorical quality is no longer apparent to speakers. It is doubtful, for example, whether anyone taking advantage of the commuter air service between London and Brussels or between New York and Washington thinks of looms or sewing machines when they talk of catching a *shuttle*. The vocabulary of a language is littered with fossilized metaphors such as these, and this continuing process makes it difficult to decide the point at which the use of a word is literal rather than figurative. Facts such as these have led some linguists, notably George Lakoff (Lakoff and Johnson 1980; Lakoff 1987), to claim that there is no principled distinction between literal and metaphorical uses of language. Such scholars see metaphor as an integral part of human categorization: a basic way of organizing our thoughts about the world. Lakoff and Johnson identify clusterings of metaphoric uses, giving them labels such as 'Time is money' to explain clusters such as 1.31 (Lakoff and Johnson 1980: 7):

1.31   You're **wasting** my time.
        This gadget will **save** you hours.
        I don't have the time to **give** you.
        How do you **spend** your time these days?
        That flat tire **cost** me an hour.
        I've **invested** a lot of time in her.

Their claim is that whole semantic fields are systematically organized around central metaphors such as these, and that their use is not just an isolated stylistic effect: that we think, culturally, of time as a commodity.

Clearly, if sentences like *How do you spend your time these days?* are identified as metaphorical, then it will prove difficult to find any uses of language that are literal. Many linguists, however, would deny that this use of *spend* is metaphorical. The position adopted by many semanticists is that this is an example of a faded or dead metaphor. The idea is that metaphors fade over time, and become part of normal literal language, much as we described for *shuttle* above. In this approach, there *is* a valid distinction between literal and non-literal language. In what we can call the **literal language theory**, metaphors and other non-literal uses of language require a different processing strategy than literal language. One view is that hearers recognize non-literal uses as semantically odd, i.e. factually nonsensical like 'eating a horse' in 1.29 earlier, but then are motivated to give them some interpretation by an assumption that speakers generally are trying to make sense. The hearer then makes inferences in order to make sense out of a non-literal utterance. Clearly some figurative expressions like *eat a horse* are quite conventionalized (i.e. well on their way to being 'dead') and do not require much working out. Other examples of non-literal language might require a little more interpretative effort, as when a reader gets to this exchange in Sean O'Faolain's novel *And Again?* (1972: 82):

1.32   'Of course,' my host said with a sigh, 'the truth is he didn't get on with the wife.'
        'Really?'
        'She flew her kite a bit too often. All Dublin knew it.'

In the literal language theory, the reader's task here is firstly to reject the literal interpretation, that the husband had a phobia about kite flying, and then to work out what kind of behaviour is being referred to so obliquely here.

We discuss hearers' assumptions about speakers' intentions in chapter 7, when we also investigate the inferences hearers routinely make to interpret utterances. In chapter 11 we discuss arguments from writers in **cognitive semantics**, like Lakoff (1987), that the literal language theory is mistaken in viewing metaphor as something extra to, and different from, ordinary literal language.

1.6.4 Semantics and pragmatics

A similarly difficult distinction is between **semantics** and **pragmatics**. These terms denote related and complementary fields of study, both concerning the transmission of meaning through language. Drawing the line between the two fields is difficult and controversial but as a preliminary we can turn to an early use of the term **pragmatics** in Charles Morris's division of semiotics:

1.33   syntax: the formal relation of signs to each other;
        semantics: the relations of signs to the objects to which the
                      signs are applicable;
        pragmatics: the relation of signs to interpreters.
                              (adapted from Morris 1938, 1955)

Narrowing signs to *linguistic* signs, this would give us a view of pragmatics as the study of the speaker/hearer's interpretation of language, as suggested by Rudolph Carnap (1942: 9, cited in Morris 1955: 218) below:

1.34   If in an investigation explicit reference is made to the speaker, or, to put it in more general terms, to the user of a language, then we assign it to the field of pragmatics. (Whether in this case reference to designata is made or not makes no difference for this classification.) If we abstract from the user of the language and analyze only the expressions and their designata, we are in the field of semantics. And if, finally, we abstract from the designata also and analyze only the relations between the expressions, we are in (logical) syntax. The whole science of language, consisting of the three parts mentioned, is called semiotic.

We might interpret this, rather crudely, as:

1.35   meaning described in relation to speakers = pragmatics
            and hearers
        meaning abstracted away from users   = semantics.

Let's investigate what this might mean, using a simple example. A speaker can utter the same sentence to a listener, e.g. *The place is closing*, and mean to use it as a simple statement, or as a warning to hurry and get that last purchase (if they're in a department store) or drink (if in a bar). It could also be an invitation or command to leave. In fact we can imagine a whole series of uses for this simple sentence, depending on the speaker's wishes and the situation the participants find themselves in. Some semanticists would claim that there is some element of meaning common to all of these uses and that this common, non-situation-specific meaning is what semantics

is concerned with. On the other hand, the range of uses a sentence can be put to, depending on context, would be the object of study for pragmatics.

One way of talking about this is to distinguish between **sentence meaning** and **speaker meaning**. This suggests that words and sentences have a meaning independently of any particular use, which meaning is then incorporated by a speaker into the particular meaning she wants to convey at any one time. In this view semantics is concerned with sentence meaning and pragmatics with speaker meaning. We can see how this distinction might be used when we consider the use of pronouns, which as we mentioned earlier are very dependent on contextual support. For example if someone says to a listener *Is he awake?* we would say that the listener has to understand two things, amongst others, to get the meaning: the first is that in English sentence meaning *he* means something like 'male entity referred to by the speaker, not the speaker and not the person spoken to', and the second is how to work out who right now the speaker is referring to by *he*. In this view knowing the first is part of semantic knowledge and working out the second is a task for one's pragmatic competence.

The advantage of such a distinction is that it might free the semanticist from having to include all kinds of knowledge in semantics. It would be the role of pragmaticists to investigate the interaction between purely linguistic knowledge and general or encyclopaedic knowledge: an issue we touched on earlier. As we shall see in chapter 7, in order to understand utterances, hearers seem to use both types of knowledge, along with knowledge about the context of the utterance and common-sense reasoning, guesses etc. A semantics/pragmatics division enables semanticists to concentrate on just the linguistic element in utterance comprehension. Pragmatics would then be the field which studies how hearers fill out the semantic structure with contextual information (for example, work out who the speaker is referring to by pronouns etc.) and make inferences which go beyond the meaning of what was said to them (for example that *I'm tired* might mean *Let's go home*).

The semantics/pragmatics distinction seems then to be a useful one. The problems with it emerge when we get down to detail: precisely which phenomena are semantic and which pragmatic? As discussed in chapters 3 and 7, much of meaning seems to depend on context: it is often difficult, for example, to identify a meaning for a word that does not depend on the context of its use. Our strategy in this book will be not to try too hard to draw a line along this putative semantics/pragmatics divide. Some theorists are sceptical of the distinction (e.g. George Lakoff 1987, Langacker 1987, 1991), while others accept it but draw the line in different places. The reader is referred to discussions in Levinson (1983) and Mey (2001) for detail. What will become clear as we proceed is that it is very difficult to shake context out of language and that the structure of sentences minutely reveals that they are designed by their speakers to be uttered in specific contexts and with desired effects. Chapter 7 is largely devoted to providing examples of these contextual aspects of meaning.

1.7   Summary

In this chapter we have taken a brief look at the task of establishing semantics as a branch of linguistics. We identified three challenges to doing this: circularity, context and the status of linguistic knowledge. We will see examples of these problems and proposed solutions as we proceed through this book. We noted that establishing a semantics component in linguistic theory involves deciding how to relate word meaning and sentence meaning. Finally, we introduced some background ideas that are assumed in many semantic theories and which we will examine in more detail in subsequent chapters: reference and sense; utterance, sentence and proposition; literal and non-literal meaning; and semantics and pragmatics. We turn to reference and sense in the next chapter.

FURTHER READING

A concise general history of linguistics is Robins (1990) and the influence of the ideas of de Saussure on modern linguistics is described in Lepschy (1982). Matthews (1993) describes American linguistics from Bloomfield to Chomsky. Two very detailed surveys of semantics, which include the topics mentioned in this chapter and others we will cover later, are Lyons (1977) and Allan (1986). These both consist of two volumes and are very useful as works of reference. An introduction to the areas covered by pragmatics is given by Mey (2001).

EXERCISES

1.1   We made the claim that meaning is **compositional**, that is, that the meaning of complex linguistic expressions is built up from the meaning of their constituent parts. However there are a number of areas where compositionality is restricted and one of these is compound words. Below is a list of English compound nouns. One very common pattern is for the second element to identify the type of thing the compound is, while the first is some kind of qualifier. So a *teacup* is a kind of cup out of which tea may be drunk. Divide the list below into two types: one where the meaning is predictable from the meaning of the two parts and a second type where the meaning is not predictable in this way. For the first type, which show a certain compositionality, how would you characterize the type of qualification made by the first part of the compound? Check your explanations against a dictionary's entries.

blackmail	foxhound	leisure centre	software
boyfriend	greenhouse	mouse mat	sunstroke
businessman	half-sister	redhead	taste bud
daydream	hotdog	six-pack	textbook
deadlock	houseboat	sky-scraper	spin-doctor
flight deck	housewife	softball	windsock

1.2   We raised the issue of a speaker's **linguistic** and **encyclopaedic knowledge**. Most English speakers will have encountered the words below, which we partly define below by their part of speech and some indication of context of use. Try to give an exact definition of their meanings, as if you were writing your own dictionary:

yew   (noun: a tree)
copper   (noun: a metal)
vodka   (noun: a drink)
hay   (noun: farming product)

How would you distinguish between the following pairs, using your original definitions as a basis?

yew/oak   copper/bronze   vodka/gin   hay/straw

When you have done this exercise, you may like to compare your definitions against a dictionary.

1.3   We used the term **reference** for the use of nominals (noun phrases and names) and pronouns to identify or pick out individuals in the world. For each of the following, imagine the sentence being spoken in an average kind of situation. Discuss which elements would be used to **refer** in your situation.

a.   This schedule is crazy.
b.   She enjoyed herself at the party.
c.   There's a policeman looking at your car.
d.   The script calls for a short fat guy.
e.   You asked for a ham sandwich; this is a ham sandwich.

1.4   Discuss the use of **figurative** language in the following extracts from (a) *The Economist* magazine and (b) *The New Scientist* magazine:

a.   The recent verdict of *Fortune* magazine was that AT&T is 'dying', and who could disagree? The telecoms giant's sales

are falling, predators are snapping at its heels, and the ambitions of its boss, Michael Armstrong, lie in expensive rubble around him. Even Golden Boy, the statue that decorates the firm's New Jersey headquarters, is up for sale: the company is seeking more modest accommodation elsewhere. If it really is dying, however, nobody seems to have told AT&T. The company thinks itself in ruddy health.[10]

b. But why should cells want to detect light? The most obvious answer is that they are talking to one another, says Albrecht-Buehler. Cells in embryos might signal with photons so that they know how and where to fit into the developing body. And now he wants to learn their language. He envisages doctors telling cells what they want them to do in words they understand. You might tell cancer cells to stop growing or encourage cells near wounds to start again. 'We may learn to compose our own messages in the language of cells to compel them to carry out specialized tasks that they've never performed.'[11]

NOTES

1 For an accessible introduction to semiotics, see Sebeok (1994). A more advanced discussion is in Eco (1976).
2 There are, however, iconic elements to language, as for example the use of **onomatopoeia**, or sound symbolism, as in the English words *tick-tock*, *cuckoo*, *ratatat* and *sizzle*. Some writers claim that iconicity is a much more extensive feature of language than this; see Haiman (1985) for example.
3 To avoid cumbersome devices like 's/he', we will when discussing simple conversations use 'he' and 'she' at random.
4 For introductions to pragmatics see Levinson (1983) and Mey (2001).
5 We look at semantics within this Cognitive Grammar approach in chapter 11.
6 As mentioned earlier, in Cognitive Grammar (Langacker 1987, 1996), discussed later in chapter 11, no distinction is made between semantic and grammatical rules.
7 This distinction between sense and reference is a translation of Frege's distinction between *Sinn* and *Bedeutung*; see Frege (1980), especially the section 'On Sense and Reference' (originally published in 1892). We discuss these notions further in chapter 2.
8 See Allwood et al. (1977) for details of translating from English sentences into such logical formulae. We will look at this strategy again in chapter 10.
9 For simplicity this section has concentrated on the relationship between propositions and the utterance of full sentences. In fact, as we can see from examples 1 and 2 below, in the right context propositions can be communicated by less than full sentences:

1 What's the longest river in the world?

2 a. The Nile is the longest river in the world.
   b. The Nile is.
   c. The Nile.

It seems reasonable to say that in the context of the question in 1 above, each of 2a–c can communicate the proposition THE NILE IS THE LONGEST RIVER IN THE WORLD, even though only 2a is a full sentence; 2b is a reduced or elliptical sentence, while 2c is of course just a noun phrase. This is another example of the possible indirectness of the relationship between utterances, sentences and propositions: a proposition can be communicated by the utterance of various grammatical units, one of which is a sentence. See Lyons (1981: 195ff.) for discussion of this point. We assume here that grammatical units like sentence (S), noun phrase (NP), verb phrase (VP) etc. are defined and specified at the level of syntax.

10 From *The Economist*, 23 February–1 March 2002: 69.
11 From *The New Scientist*, 23 February 2002: 33.